JUNIPER NETWORKS | Engineering Simplicity

Junos® OS

Traffic Sampling, Forwarding, and Monitoring User Guide

Published
2023-03-19

JUNOS

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Junos® OS Traffic Sampling, Forwarding, and Monitoring User Guide
Copyright © 2023 Juniper Networks, Inc. All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at https://support.juniper.net/support/eula/. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About This Guide | xi

1 Overview

Traffic Sampling, Forwarding, and Monitoring Overview | 2

2 Collecting Traffic Samples for Network Monitoring

Traffic Sampling Configuration | 4

Minimum Traffic Sampling Configuration | 5

Configuring Traffic Sampling | 6

Disabling Traffic Sampling | 9

Collecting Traffic Sampling Output in a File | 10

Directing Traffic Sampling Output to a Server Running the cflowd Application | 12

Collecting Traffic Sampling Output in the Cisco Systems NetFlow Services Export Version 9 Format | 16

Example: Sampling a Single SONET/SDH Interface | 18

Example: Sampling All Traffic from a Single IP Address | 19

Example: Sampling All FTP Traffic | 21

Tracing Traffic-Sampling Operations | 22

3 Configuring Traffic Forwarding for Network Monitoring

Configuring Traffic Forwarding and Monitoring | 24

Configuring IPv4 and IPv6 Accounting | 29

Configuring Discard Accounting | 31

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers | 32

Configuring Passive Flow Monitoring | 35

Configuring Port Mirroring | 37

Configuring Next-Hop Groups to Use Multiple Interfaces to Forward Packets Used in Port Mirroring | 42

Defining a Port-Mirroring Firewall Filter | 44

Defining a Next-Hop Group on MX Series Routers for Port Mirroring | 47

4 Configuring Forwarding Table Filters to Efficiently Route Traffic

Configuring Forwarding Table Filters | 51

Forwarding Table Filters for Routing Instances on ACX Series Routers | 53

Applying Forwarding Table Filters | 54

5 Configuring Forwarding Options for Load Balancing Traffic

Configuring Load Balancing for Ethernet Pseudowires | 59

Configuring Load-Balance Groups | 60

Understanding the Algorithm Used to Load Balance Traffic on MX Series Routers | 61

Understanding Per-Packet Load Balancing | 75

Configuring Per-Packet Load Balancing | 77

Understanding Load Balancing for BGP Traffic with Unequal Bandwidth Allocated to the Paths | 81

Understanding the Default BGP Routing Policy on Packet Transport Routers (PTX Series) | 82

Per-Flow and Per-Prefix Load Balancing Overview | 85

Configuring Per-Prefix Load Balancing | 85

Configuring Per-Flow Load Balancing Based on Hash Values | 87

Configuring Load Balancing Based on MAC Addresses | 88

Load Balancing VPLS Non-Unicast Traffic Across Member Links of an Aggregate Interface | 89

Example: Configuring Multicast Load Balancing over Aggregated Ethernet Links | 90

Requirements | 91

Overview | 91

Configuration | 92

Verification | 107

6 Configuring Other Forwarding Options

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents | 109

Configuring DNS and TFTP Packet Forwarding | 112

Configuring Port-based LAN Broadcast Packet Forwarding | 117

Preventing DHCP Spoofing on MX Series 5G Universal Routing Platforms | 119

Understanding the Hyper Mode Feature on Enhanced MPCs for MX Series Routers and EX9200 Switches | 121

Configuring Hyper Mode on Enhanced MPCs to Speed Up Packet Processing | 124

Unsupported Features and CLI Commands When Hyper Mode Is Enabled | 126

7 Configuration Statements

accounting | 133

aggregation | 135

autonomous-system-type | 137

bootp | 139

bum-hashing | 141

cflowd (Discard Accounting) | 142

cflowd (Flow Monitoring) | 144

client-address | 146

client-response-ttl | 148

description (Forwarding Options) | 150

dhcp-relay (DHCP Spoofing Prevention) | 151

disable (Forwarding Options) | 153

domain | 155

ecmp-local-bias | 156

enhanced-hash-key | 158

export-format | 164

family (Filtering) | 166

family (Monitoring) | 168

family (Port Mirroring) | 169

family (Sampling) | 172

family inet | 174

family mpls | 177

family multiservice | 181

file (Extended DHCP Relay Agent and Helpers Trace Options) | 184

file (Sampling) | 186

file (Trace Options) | 187

filename (Sampling) | 189

files (Sampling and Traceoptions) | 190

filter (IPv4, IPv6, and MPLS) | 192

filter (VPLS) | 194

flood | 195

flow-active-timeout | 197

flow-export-destination | 199

flow-inactive-timeout | 200

flow-server | 202

group (DHCP Spoofing Prevention) | 204

gtp-tunnel-endpoint-identifier | 206

hash-key (Forwarding Options) | 208

helpers | 212

hosted-service-identifier | 216

hosted-services | 218

hyper-mode (forwarding-options) | 219

indexed-load-balance | 221

input (Forwarding Table) | 223

input (Port Mirroring) | 225

input (Sampling) | 226

instance | 228

interface (Accounting or Sampling) | 230

interface (BOOTP) | 232

interface (DHCP Spoofing Prevention) | 234

interface (DNS, Port, and TFTP Packet Forwarding or Relay Agent) | 235

interface (Monitoring) | 237

interface (Next-Hop Group) | 239

interface (Port Mirroring) | 241

l2tp-tunnel-session-identifier | 243

link-layer-broadcast-inet-check | 244

load-balance (Forwarding Options) | 246

load-balance-group | 249

local-bias | 250

local-dump | 252

max-packets-per-second | 253

maximum-hop-count | 255

maximum-packet-length | 257

minimum-wait-time | 259

mirror-once | 261

monitoring | 262

next-hop (Forwarding Options) | 264

next-hop-group (Forwarding Options) | 266

next-hop-group | 268

no-filter-check | 270

no-listen | 272

output (Accounting) | 273

output (Forwarding Table) | 275

output (Monitoring) | 277

output (Port Mirroring) | 278

output (Sampling) | 281

per-flow | 283

per-prefix | 285

port (cflowd) | 286

port (Packet Forwarding) | 288

port-mirroring | 291

rate (Forwarding Options) | 296

relay-agent-option | 298

route-accounting | 300

run-length | 302

sampling (Forwarding Options) | 304

server (DHCP and BOOTP Relay Agent) | 308

server (DNS, Port, and TFTP Service) | 310

server-address (Hosted Services) | 312

server-profile | 313

server-profile (Active Flow Monitoring) | 315

size (Sampling and Traceoptions) | 316

source-checking | 318

stamp | 320

tftp | 321

traceoptions (DNS, Port, and TFTP Packet Forwarding) | 323

traceoptions (Port Mirroring and Traffic Sampling) | 326

vector routing | 328

vector-routing | 328

8 traceoptions (Port Mirroring and Traffic Sampling)

transport-type | 334

version | 335

version9 | 337

world-readable (Forwarding Options) | 338

9 Operational Commands

clear passive-monitoring statistics | 343

clear services flow-collector statistics | 344

request services flow-collector change-destination primary interface | 346

request services flow-collector change-destination secondary interface | 348

request services flow-collector test-file-transfer | 350

show chassis forwarding | 352

show forwarding-options hyper-mode | 355

show forwarding-options load-balance | 357

show forwarding-options port-mirroring | 362

show forwarding-options next-hop-group | 365

show interfaces (Flow Monitoring) | 370

show interfaces (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet) | 378

show interfaces statistics | 403

show passive-monitoring error | 422

show passive-monitoring flow | 426

show passive-monitoring memory | 429

show passive-monitoring status | 432

show passive-monitoring usage | 435

show route forwarding-table | 438

show services accounting aggregation | 453

show services accounting aggregation template | 459

show services accounting errors | 462

show services accounting flow | 469

show services accounting flow-detail | 478

show services accounting memory | 485

show services accounting packet-size-distribution | 488

show services accounting status | 491

show services accounting usage | 497

show services flow-collector file interface | 500

show services flow-collector input interface | 504

show services flow-collector interface | 507

About This Guide

This guide provides information about traffic sampling, which allows you to sample IP traffic based on particular input interfaces and various fields in the packet header. You can also use traffic sampling to monitor interfaces, protocols, and addresses.

The guide also offers information about per-flow load balancing, port mirroring, and domain name system (DNS) or Trivial File Transfer Protocol (TFTP) forwarding.

Traffic sampling and forwarding are supported only on routers equipped with an Internet Processor II application-specific integrated circuit (ASIC). To determine whether a routing platform has an Internet Processor II ASIC, use the `show chassis hardware` command.

1
CHAPTER

Overview

Traffic Sampling, Forwarding, and Monitoring Overview | 2

Traffic Sampling, Forwarding, and Monitoring Overview

Traffic sampling allows you to sample IP traffic based on particular input interfaces and various fields in the packet header. You can also use traffic sampling to monitor any combination of specific logical interfaces, specific protocols on one or more interfaces, a range of addresses on a *logical interface*, or individual IP addresses. Information about the sampled packets is saved to files on the router's hard disk.

Traffic sampling is not meant to capture all packets received by a router. We do not recommend excessive sampling (a rate greater than 1/1000 packets), because it can increase the load on your processor. If you need to set a higher sampling rate to diagnose a particular problem or type of traffic received, we recommend that you revert to a lower sampling rate after you discover the problem or troublesome traffic. In addition, traffic sampling and forwarding are supported only on routers equipped with an Internet Processor II application-specific integrated circuit (ASIC). To determine whether a routing platform has an Internet Processor II ASIC, use the `show chassis hardware` command.

Junos OS supports both per-packet and per-flow load balancing. In Junos OS Release 9.0 and later, you can configure per-prefix load balancing. This feature enables the router to elect the next hop independent of the route chosen by other routers. The result is a better utilization of available links. Likewise, you can configure Junos OS so that, for the active route, all next-hop addresses for a destination are installed in the forwarding table. This is called per-packet load balancing, which you can use to spread traffic across multiple paths between routers.

With forwarding policies, you can configure per-flow load balancing, *port mirroring*, and domain name system (DNS) or Trivial File Transfer Protocol (TFTP) forwarding.

Release History Table

Release	Description
9.0	In Junos OS Release 9.0 and later, you can configure per-prefix load balancing.

2

CHAPTER

Collecting Traffic Samples for Network Monitoring

Traffic Sampling Configuration | 4

Minimum Traffic Sampling Configuration | 5

Configuring Traffic Sampling | 6

Disabling Traffic Sampling | 9

Collecting Traffic Sampling Output in a File | 10

Directing Traffic Sampling Output to a Server Running the cflowd Application | 12

Collecting Traffic Sampling Output in the Cisco Systems NetFlow Services Export Version 9 Format | 16

Example: Sampling a Single SONET/SDH Interface | 18

Example: Sampling All Traffic from a Single IP Address | 19

Example: Sampling All FTP Traffic | 21

Tracing Traffic-Sampling Operations | 22

Traffic Sampling Configuration

To configure traffic sampling, include the `sampling` statement at the `[edit forwarding-options]` hierarchy level:

```
[edit forwarding-options]
sampling {
    disable;
    family (inet | inet6 | mpls) {
        disable;
        output {
            aggregate-export-interval seconds;
            extension-service service-name;
            file {
                disable;
                filename filename;
                files number;
                size bytes;
                (stamp | no-stamp);
                (world-readable | no-world-readable);
            }
            flow-active-timeout seconds;
            flow-inactive-timeout seconds;
            flow-server hostname {
                aggregation {
                    autonomous-system;
                    destination-prefix;
                    protocol-port;
                    source-destination-prefix {
                        caida-compliant;
                    }
                    source-prefix;
                }
                autonomous-system-type (origin | peer);
                (local-dump | no-local-dump);
                port port-number;
                source-address address;
                version format;
                version9 {
                    template template-name;
```

```
                }
            }
            interface interface-name {
                engine-id number;
                engine-type number;
                source-address address;
            }
        }
    }
    input {
        max-packets-per-second number;
        maximum-packet-length bytes;
        rate number;
        run-length number;
    }
    traceoptions {
        file filename {
            files number;
            size bytes;
            (world-readable | no-world-readable);
        }
    }
}
```

Minimum Traffic Sampling Configuration

To configure traffic sampling, you must perform at least the following tasks:

1. Create a firewall filter to apply to the logical interfaces being sampled by including the `filter` statement at the `[edit firewall family family-name]` hierarchy level. In the filter `then` statement, you must specify the action modifier **sample** and the action **accept**.

```
[edit firewall family family-name]
filter filter-name {
    term term-name {
        then {
            sample;
            accept;
        }
```

```
    }
}
```

2. Apply the filter to the interfaces on which you want to sample traffic:

```
[edit interfaces]
interface-name {
    unit logical-unit-number {
        family family-name {
            filter {
                input filter-name;
            }
            address address {
                destination destination-address;
            }
        }
    }
}
```

3. Enable sampling and specify a nonzero sampling rate:

```
[edit forwarding-options]
sampling {
    input {
        rate number;
    }
}
```

Configuring Traffic Sampling

On routing platforms containing a Monitoring Services PIC or an Adaptive Services PIC, you can configure traffic sampling for traffic passing through the routing platform. In Junos OS Release 8.3 and later, you can also configure traffic sampling of MPLS traffic.

To configure traffic sampling on a logical interface:

1. Include the `input` statement at the `[edit forwarding-options sampling]` hierarchy level, for example:

```
[edit forwarding-options sampling]
input {
    max-packets-per-second number;
    maximum-packet-length bytes
    rate number;
    run-length number;
}
```

Junos OS Release 17.2R1, you can export flow records generated by inline flow monitoring to four collectors under a family with the same source IP address. The Packet Forwarding Engine (PFE) can export the flow record, flow record template, option data, and, option data template packet to all configured collectors. You can configure the multiple collectors at the `[edit forwarding-options sampling instance instance name]` hierarchy level.

NOTE: You cannot change the source IP address for collectors under the same family.

2. Specify the threshold traffic value by using the `max-packets-per-second` statement. The value is the maximum number of packets to be sampled, beyond which the sampling mechanism begins dropping packets. The range is 0 through 65,535. A value of 0 instructs the Packet Forwarding Engine not to sample any packets. The default value is 1000.

NOTE: This statement is not valid for port mirroring.

3. Specify the maximum length of the sampled packet by using the `maximum-packet-length bytes` statement. For *bytes*, specify a value.

NOTE: For MX-Series devices with Modular Port Concentrators (MPCs) and T4000 router with Type 5 FPC, port-mirrored or sampled packets can be truncated (or clipped) to any length in the range of 1 to 255 bytes. Only 1 to 255 are valid values for packet truncation on these devices. For other devices, the range is from 0 to 9216. A maximum-packet-length value of zero represents that truncation is disabled, and the entire packet is mirrored or sampled.

4. Specify the sampling rate by setting the values for *rate* and *run-length* (see Figure 1 on page 8).

Figure 1: Configure Sampling Rate

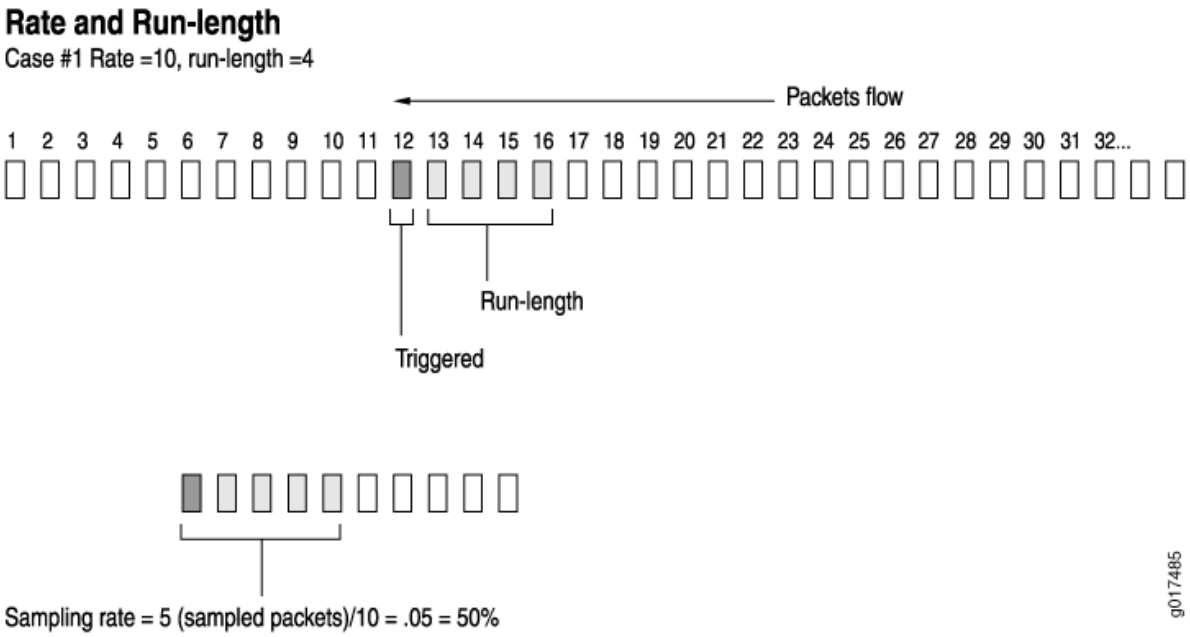

The forwarding plane provides support for random sampling that can be configured through the *rate* or `run-length` statement. The `rate` statement sets the ratio of the number of packets to be sampled on an average. For example, if you configure a rate of 10, on average every tenth packet (1 packet out of 10) is sampled.

The `run-length` statement specifies the number of matching packets to sample following the initial one-packet trigger event. Configuring a run length greater than 0 allows you to sample packets following those already being sampled.

NOTE: The `run-length` statement is not supported on MX Series routers with Modular Port Concentrators (MPCs) and T4000 router with Type 5 FPC.

You can also send the sampled packets to a specified host using the cflowd version 5 and 8 formats or the version 9 format as defined in RFC 3954. For more information, see "Directing Traffic Sampling Output to a Server Running the cflowd Application" on page 12 and "Collecting Traffic Sampling Output in the Cisco Systems NetFlow Services Export Version 9 Format" on page 16.

Junos OS does not sample packets originating from the router. If you configure a sampling filter and apply it to the output side of an interface, then only the transit packets going through that interface are sampled. Packets that are sent from the Routing Engine to the Packet Forwarding Engine are not sampled.

When you apply a firewall filter to a loopback interface, the filter might block responses from the Monitoring Services PIC. To allow responses from the Monitoring Services PIC to pass through for

sampling purposes, configure a term in the firewall filter to include the Monitoring Services PIC's IP address.

NOTE: Targeted broadcast does not work when the targeted broadcast option `forward-and-send-to-re` and the traffic sampling option `sampling` are configured on the same egress interface of an M320 router, a T640 router, or an MX960 router. To overcome this scenario, you must either disable one of the these options or enable the `sampling` option with the targeted broadcast option `forward-only` on the egress interface. For information about targeted broadcast, see *Understanding Targeted Broadcast*.

RELATED DOCUMENTATION

Guidelines for Configuring Firewall Filters

Guidelines for Applying Standard Firewall Filters

Disabling Traffic Sampling

To explicitly disable traffic sampling on the router, include the `disable` statement at the `[edit forwarding-options sampling]` hierarchy level:

```
[edit forwarding-options sampling]
disable;
```

NOTE: The `disable` statement at the `[edit forwarding-options sampling]` hierarchy level disables only Routing Engine-based sampling. To disable PIC-based sampling and inline sampling, include the `disable` statement at the `[edit forwarding-options sampling instance instance-name]` hierarchy level.

Collecting Traffic Sampling Output in a File

IN THIS SECTION

- Traffic Sampling Output Format | 11

You configure traffic sampling results to a file in the **/var/tmp** directory. To collect the sampled packets in a file, include the `file` statement at the `[edit forwarding-options sampling output]` hierarchy level:

```
[edit forwarding-options sampling family family-name output]
file <disable> filename filename <files number> <size bytes> <stamp | no-stamp > <world-readable |
no-world-readable>;
```

To configure the period of time before an active flow is exported, include the `flow-active-timeout` statement at the `[edit forwarding-options sampling output family (inet | inet6 | mpls)]` hierarchy level:

```
[edit forwarding-options sampling family (inet | inet6 | mpls) output]
flow-active-timeout seconds;
```

To configure the period of time before a flow is considered inactive, include the `flow-inactive-timeout` statement at the `[edit forwarding-options sampling output]` hierarchy level:

```
[edit forwarding-options sampling family (inet | inet6 | mpls) output]
flow-inactive-timeout seconds;
```

To configure the interface that sends out monitored information, include the `interface` statement at the `[edit forwarding-options sampling output]` hierarchy level:

```
[edit forwarding-options sampling family (inet | inet6 | mpls) output]
interface interface-name {
    engine-id number;
    engine-type number;
    source-address address;
}
```

NOTE: This feature is not supported with the version 9 template format. You must send traffic flows collected using version 9 to a server. For more information see "Collecting Traffic Sampling Output in the Cisco Systems NetFlow Services Export Version 9 Format" on page 16.

Traffic Sampling Output Format

Traffic sampling output is saved to an ASCII text file. The following is an example of the traffic sampling output that is saved to a file in the **/var/tmp** directory. Each line in the output file contains information for one sampled packet. You can optionally display a timestamp for each line.

The column headers are repeated after each group of 1000 packets.

```
# Apr  7 15:48:50
Time                    Dest           Src Dest Src Proto TOS Pkt Intf  IP   TCP
                        addr          addr port port          len num frag flags
Apr 7 15:48:54 192.168.9.194 192.168.9.195   0    0   1   0x0  84  8   0x0   0x0
Apr 7 15:48:55 192.168.9.194 192.168.9.195   0    0   1   0x0  84  8   0x0   0x0
Apr 7 15:48:56 192.168.9.194 192.168.9.195   0    0   1   0x0  84  8   0x0   0x0
Apr 7 15:48:57 192.168.9.194 192.168.9.195   0    0   1   0x0  84  8   0x0   0x0
Apr 7 15:48:58 192.168.9.194 192.168.9.195   0    0   1   0x0  84  8   0x0   0x0
```

The output contains the following fields:

- **Time**—Time at which the packet was received (displayed only if you include the `stamp` statement in the configuration)
- **Dest addr**—Destination IP address in the packet
- **Src addr**—Source IP address in the packet
- **Dest port**—Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) port for the destination address
- **Src port**—TCP or UDP port for the source address
- **Proto**—Packet's protocol type
- **TOS**—Contents of the type-of-service (ToS) field in the IP header
- **Pkt len**—Length of the sampled packet, in bytes

- **Intf num**—Unique number that identifies the sampled logical interface
- **IP frag**—IP fragment number, if applicable
- **TCP flags**—Any TCP flags found in the IP header

To set the timestamp option for the file **my-sample**, enter the following:

```
[edit forwarding-options sampling family (inet | inet6 | mpls) output file]
user@host# set filename my-sample files 5 size 2m world-readable stamp;
```

Whenever you toggle the timestamp option, a new header is included in the file. If you set the **stamp** option, the **Time** field is displayed.

```
# Apr  7 15:48:50
# Time            Dest        Src  Dest   Src Proto  TOS   Pkt  Intf    IP   TCP
#                 addr       addr  port  port             len   num  frag flags
# Feb  1 20:31:21
#                 Dest        Src  Dest   Src Proto  TOS   Pkt  Intf    IP   TCP
#                 addr       addr  port  port             len   num  frag flags
```

Directing Traffic Sampling Output to a Server Running the cflowd Application

IN THIS SECTION

- Debugging cflowd Flow Aggregation | **15**

You can collect an aggregate of sampled flows and send the aggregate to a specified host that runs the cflowd application available from the Cooperative Association for Internet Data Analysis (CAIDA) (**http://www.caida.org**). By using cflowd, you can obtain various types of byte and packet counts of flows through a router.

The cflowd application collects the sampled flows over a period of 1 minute. At the end of the minute, the number of samples to be exported are divided over the period of another minute and are exported over the course of the same minute.

Before you can perform flow aggregation, the routing protocol process must export the autonomous system (AS) path and routing information to the sampling process. To do this, include the `route-record` statement:

```
route-record;
```

You can include this statement at the following hierarchy levels:

- `[edit routing-options]`
- `[edit routing-instances` *routing-instance-name* `routing-options]`

By default, flow aggregation is disabled. To enable the collection of flow aggregates, include the `flow-server` statement at the `[edit forwarding-options sampling output]` hierarchy level:

```
[edit forwarding-options sampling family (inet | inet6 | mpls) output ]
flow-server hostname {
    aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
            caida-compliant;
        }
        source-prefix;
    }
    autonomous-system-type (origin | peer);
    (local-dump | no-local-dump);
    port port-number;
    source-address address;
    version format;
}
```

In the cflowd statement, specify the name, identifier, and source-address of the host that collects the flow aggregates. You must also include the UDP port number on the host and the **version**, which gives the format of the exported cflowd aggregates. To specify an IPv4 source address, include the `source-address` statement. To collect cflowd records in a log file before exporting, include the `local-dump` statement. To specify the cflowd version number, include the `version` statement. The cflowd version is either 5 or 8.

You can specify both host (cflowd) sampling and port mirroring in the same configuration. You can perform RE-sampling and port mirroring actions simultaneously. However, you cannot perform PIC-sampling and port mirroring actions simultaneously.

To specify aggregation of specific types of traffic, include the `aggregation` statement. This conserves memory and bandwidth enabling cflowd to export targeted flows rather than all the aggregated

NOTE: Aggregation is valid only if cflowd version 8 is specified.

To specify a flow type, include the `aggregation` statement at the `[edit forwarding-options sampling output cflowd` *`hostname`*`]` hierarchy level:

```
[edit forwarding-options sampling family (inet | inet6 | mpls) output hostname]
aggregation {
    source-destination-prefix;
}
```

You specify the aggregation type using one of the following options:

- **autonomous-system**—Aggregate by AS number; may require setting the separate cflowd `autonomous-system-type` statement to include either **origin** or **peer** AS numbers. The **origin** option specifies to use the origin AS of the packet source address in the Source Autonomous System cflowd field. The **peer** option specifies to use the peer AS through which the packet passed in the Source Autonomous System cflowd field. By default, **cflowd** exports the origin AS number.
- **destination-prefix**—Aggregate by destination prefix (only).
- **protocol-port**—Aggregate by protocol and port number; requires setting the separate `cflowd port` statement.
- **source-destination-prefix**—Aggregate by source and destination prefix. Version 2.1b1 of CAIDA's cflowd application does not record source and destination mask length values in compliance with CAIDA's *cflowd Configuration Guide*, dated August 30, 1999. If you configure the `caida-compliant` statement, Junos OS complies with Version 2.1b1 of cflowd. If you do not include the `caida-compliant` statement in the configuration, Junos OS records source and destination mask length values in compliance with the *cflowd Configuration Guide*.
- **source-prefix**—Aggregate by source prefix (only).

Collection of sampled packets in a local ASCII file is not affected by the `cflowd` statement.

Debugging cflowd Flow Aggregation

To collect the cflowd flows in a log file before they are exported, include the **local-dump** option at the `[edit forwarding-options sampling output cflowd hostname]` hierarchy level:

```
[edit forwarding-options sampling family (inet | inet6 | mpls) output flow-server hostname]
local-dump;
```

By default, the flows are collected in **/var/log/sampled**; to change the filename, include the `filename` statement at the `[edit forwarding-options sampling traceoptions]` hierarchy level. For more information about changing the filename, see "Collecting Traffic Sampling Output in a File" on page 10.

> **NOTE**: Because the **local-dump** option adds extra overhead, you should use it only while debugging cflowd problems, not during normal operation.

The following is an example of the flow information. The AS number exported is the origin AS number. All flows that belong under a cflowd header are dumped, followed by the header itself:

```
Jun 27 18:35:43 v5 flow entry
Jun 27 18:35:43    Src addr: 10.53.127.1
Jun 27 18:35:43    Dst addr: 10.6.255.15
Jun 27 18:35:43    Nhop addr: 192.168.255.240
Jun 27 18:35:43    Input interface: 5
Jun 27 18:35:43    Output interface: 3
Jun 27 18:35:43    Pkts in flow: 15
Jun 27 18:35:43    Bytes in flow: 600
Jun 27 18:35:43    Start time of flow: 7230
Jun 27 18:35:43    End time of flow: 7271
Jun 27 18:35:43    Src port: 26629
Jun 27 18:35:43    Dst port: 179
Jun 27 18:35:43    TCP flags: 0x10
Jun 27 18:35:43    IP proto num: 6
Jun 27 18:35:43    TOS: 0xc0
Jun 27 18:35:43    Src AS: 64496
Jun 27 18:35:43    Dst AS: 64511
Jun 27 18:35:43    Src netmask len: 16
Jun 27 18:35:43    Dst netmask len: 0
```

[... 41 more **v5 flow** entries; then the following header:]

```
Jun 27 18:35:43 cflowd header:
Jun 27 18:35:43   Num-records: 42
Jun 27 18:35:43   Version: 5
Jun 27 18:35:43   Flow seq num: 118
Jun 27 18:35:43   Engine id: 0
Jun 27 18:35:43   Engine type: 3
```

Collecting Traffic Sampling Output in the Cisco Systems NetFlow Services Export Version 9 Format

IN THIS SECTION

- Example: Configuring Active Flow Monitoring Using Version 9 | **17**

In Junos OS Release 8.3 and later, you can collect a record of sampled flows using the version 9 format as defined in RFC 3954, *Cisco Systems NetFlow Services Export Version 9*. Version 9 uses templates to collect a set of sampled flows and send the record to a specified host.

You configure the version 9 template used to collect a record of sampled flows at the `[edit services monitoring]` hierarchy level. For more information, see the Junos OS Services Interfaces Library for Routing Devices and the Monitoring, Sampling, and Collection Services Interfaces User Guide.

To enable the collection of traffic flows using the version 9 format, include the `version9` statement at the `[edit forwarding-options sampling family` *`family-name`* `output flow-server` *`hostname`*`]` hierarchy level:

```
[edit forwarding-options sampling family family-name output flow-server hostname]
version9 {
    template template-name;
}
```

template-name is the name of the version 9 template configured at the `[edit services monitoring]` hierarchy level.

You configure traffic sampling at the `[edit forwarding-options sampling input]` hierarchy level. In Junos OS Release 8.3 and later, you can configure sampling for MPLS traffic as well as IPv4 traffic. You can define a version 9 flow record template suitable for IPv4 traffic, MPLS traffic, or a combination of the two. In Junos OS Release 9.5 and later, you can sample packets from both the **inet** and **mpls** protocol families at the same time. In Junos OS Release 10.4 and later, you can configure sampling for peer AS billing traffic for the **inet** and **ipv6** protocols only. For more information about how to configure traffic sampling, see "Configuring Traffic Sampling" on page 6.

The following restrictions apply to configuration of the version 9 format:

- You can configure only one host to collect traffic flows using the version 9 format. Configure the host at the `[edit forwarding-options sampling family family-name output flow-server hostname]` hierarchy level.
- You cannot specify both the version 9 format and cflowd versions 5 and 8 formats in the same configuration. For more information about how to configure flow monitoring using cflowd version 8, see "Directing Traffic Sampling Output to a Server Running the cflowd Application" on page 12.
- Any values for **flow-active-timeout** and **flow-inactive-timeout** that you configure at the `[edit forwarding-options sampling output]` hierarchy level are overridden by the values configured in the version 9 template.
- Version 9 does not support Routing Engine-based sampling. You cannot configure version 9 to send traffic sampling result to a file in the **/var/tmp** directory.

Example: Configuring Active Flow Monitoring Using Version 9

In this example, you enable active flow monitoring using version 9. You specify a template **mpls** that you configure at the `[edit services monitoring]` hierarchy level. You also configure the traffic family **mpls** to sample MPLS packets.

```
[edit forwarding-options]
sampling {
    input {
        rate 1;
        run-length;
    }
    family inet {
        output {
            flow-server 10.60.2.1 { # The IP address and port of the host
                port 2055; # that collects the sampled traffic flows.
                source-address 3.3.3.1;
                version9 {
```

```
                    template mpls; # Version 9 records are sent
                } # using the template named mpls
            }
        }
    }
}
```

Example: Sampling a Single SONET/SDH Interface

The following configuration gathers statistical sampling information from a small percentage of all traffic on a single SONET/SDH interface and collects it in a file named **sonet-samples.txt**.

Create the filter:

```
[edit firewall family inet]
filter {
    sample-sonet {
        then {
            sample;
            accept;
        }
    }
}
```

Apply the filter to the SONET/SDH interface:

```
[edit interfaces]
so-0/0/1 {
    unit 0 {
        family inet {
            filter {
                input sample-sonet;
            }
            address 10.127.68.254/32 {
                destination 10.127.74.7;
            }
        }
```

```
    }
}
```

Finally, configure traffic sampling:

```
[edit forwarding-options]
sampling {
    input {
        rate 100;
        run-length 2;
    }
    family inet {
        output {
            file {
                filename sonet-samples.txt;
                files 40;
                size 5m;
            }
        }
}
```

Example: Sampling All Traffic from a Single IP Address

The following configuration gathers statistical information about every packet entering the router on a specific Gigabit Ethernet port originating from a single source IP address of **10.45.92.31**, and collects it in a file named **samples-10-45-92-31.txt**.

Create the filter:

```
[edit firewall family inet]
filter one-ip {
    term get-ip {
        from {
            source-address 10.45.92.31;
        }
        then {
```

```
            sample;
            accept;
        }
    }
}
```

Apply the filter to the Gigabit Ethernet interface:

```
[edit interfaces]
ge-4/1/1 {
    unit 0 {
        family inet {
            filter {
                input one-ip;
            }
            address 10.45.92.254;
        }
    }
}
```

Finally, gather statistics on all the candidate samples; in this case, gather all statistics:

```
[edit forwarding-options]
sampling {
    input {
        rate 1;
    }
    family inet {
        output {
            file {
                filename samples-215-45-92-31.txt;
                files 100;
                size 100k;
            }
        }
    }
}
```

Example: Sampling All FTP Traffic

The following configuration gathers statistical information about a moderate percentage of packets using FTP in the output path of a specific T3 interface, and collects the information in a file named **t3-ftp-traffic.txt**.

Create a filter:

```
[edit firewall family inet]
filter ftp-stats {
    term ftp-usage {
        from {
            destination-port [ftp ftp-data];
        }
        then {
            sample;
            accept;
        }
    }
}
```

Apply the filter to the T3 interface:

```
[edit interfaces]
t3-7/0/2 {
    unit 0 {
        family inet {
            filter {
                input ftp-stats;
            }
            address 10.35.78.254/32 {
                destination 10.35.78.4;
            }
        }
    }
}
```

Finally, gather statistics on 10 percent of the candidate samples:

```
[edit forwarding-options]
sampling {
    input {
        rate 10;
    }
    family inet {
        output {
            file {
                filename t3-ftp-traffic.txt;
                files 50;
                size 1m;
            }
        }
    }
}
```

Tracing Traffic-Sampling Operations

Tracing operations track all traffic-sampling operations and record them in a log file in the **/var/log** directory. By default, this file is named **/var/log/sampled**. The default file size is 128 KB, and 10 files are created before the first one gets overwritten.

To trace traffic-sampling operations, include the `traceoptions` statement at the `[edit forwarding-options sampling]` hierarchy level:

```
[edit forwarding-options sampling]
traceoptions {
    file <filename>  <files number> <size bytes> <world-readable | no-world-readable>;
    no-remote-trace;
}
```

CHAPTER 3

Configuring Traffic Forwarding for Network Monitoring

Configuring Traffic Forwarding and Monitoring | 24

Configuring IPv4 and IPv6 Accounting | 29

Configuring Discard Accounting | 31

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers | 32

Configuring Passive Flow Monitoring | 35

Configuring Port Mirroring | 37

Configuring Next-Hop Groups to Use Multiple Interfaces to Forward Packets Used in Port Mirroring | 42

Defining a Port-Mirroring Firewall Filter | 44

Defining a Next-Hop Group on MX Series Routers for Port Mirroring | 47

Configuring Traffic Forwarding and Monitoring

To configure forwarding options and traffic monitoring, include statements at the `[edit forwarding-options]` hierarchy level:

```
[edit forwarding-options]
accounting group-name {
    output {
        cflowd [ hostnames ] {
            aggregation {
                autonomous-system;
                destination-prefix;
                protocol-port;
                source-destination-prefix {
                    caida-compliant;
                }
                source-prefix;
            }
            autonomous-system-type (origin | peer);
            port port-number;
            version format;
        }
        flow-active-timeout seconds;
        flow-inactive-timeout seconds;
        interface interface-name {
            engine-id number;
            engine-type number;
            source-address address;
        }
    }
}
enhanced-hash-key {
    family inet {
        gtp-tunnel-endpoint-identifier;
        incoming-interface-index;
        no-destination-port;
        no-source-port;
        type-of-service;
    }
    family inet6 {
```

```
        gtp-tunnel-endpoint-identifier;
        incoming-interface-index;
        no-destination-port;
        no-source-port;
        traffic-class;
    }
    family mpls {
        incoming-interface-index;
        label-1-exp;
        no-payload;
    }
    family multiservice {
        incoming-interface-index;
        no-payload;
        outer-priority;
    }
    services-loadbalancing {
        family inet layer-3-services {
            incoming-interface-index;
            source-address;
        }
    }
}
family family-name {
    filter {
        input filter-name;
        output filter-name;
    }
    route-accounting;
}
flood {
    input filter-name;
}
hash-key {
    family inet {
        layer-3;
        layer-4;
    }
    family mpls {
        no-interface-index;
        label-1;
        label-2;
        label-3;
```

```
        no-labels;
        no-label-1-exp;
        payload {
            ether-pseudowire;
            ip {
                layer-3-only;
                port-data {
                    source-msb;
                    source-lsb;
                    destination-msb;
                    destination-lsb;
                }
            }
        }
    }
    family multiservice }
        destination-mac;
        label-1;
        label-2;
        payload {
            ip {
                layer-3-only;
            }
        }
        source-mac;
    }
}
helpers {
    bootp {
        client-response-ttl;
        description text-description;
        interface interface-group {
            client-response-ttl number;
            description text-description;
            maximum-hop-count number;
            minimum-wait-time seconds;
            no-listen;
            server address {
                logical-system logical-system-name <routing-instance [ <default> routing-
instance-names ]>;
                routing-instance [ <default> routing-instance-names ];
            }
        }
```

```
        maximum-hop-count number;
        minimum-wait-time seconds;
        relay-agent-option;
        server [ addresses ];
    }
    domain {
        description text-description;
        server < [ routing-instance routing-instance-names ] >;
        interface interface-name {
            description text-description;
            no-listen;
            server < [ routing-instance routing-instance-names ] >;
        }
    }
    tftp {
        description text-description;
        server < [ routing-instance routing-instance-names ] >;
        interface interface-name {
            description text-description;
            no-listen;
            server < [ routing-instance routing-instance-names ] >;
        }
    }
    traceoptions {
        file <filename> <files number> <match regular-expression> <size size> <world-readable |
no-world readable>;
        flag flag;
        level severity-level;
        no-remote-trace;
    }
}
load-balance {
    indexed-load-balance;
    per-flow {
        hash-seed number;
    }
    per-prefix {
        hash-seed number;
    }
}
monitoring group-name {
    family inet {
        output {
```

```
            cflowd hostname {
                port port-number;
            }
            export-format cflowd-version-5;
            flow-active-timeout seconds;
            flow-export-destination {
                cflowd-collector;
            }
            flow-inactive-timeout seconds;
             interfaceinterface-name {
                engine-id number;
                engine-type number;
                input-interface-index number;
                output-interface-index number;
                source-address address;
            }
        }
    }
}
next-hop-group [ group-names ] {
    interface interface-name {
        next-hop [ addresses ];
    }
}
port-mirroring {
    family (ccc | inet | inet6 | vpls) {
        output {
            interface interface-name {
                next-hop address;
            }
            no-filter-check;
        }
        input {
            maximum-packet-length bytes;
            rate number;
            run-length number;
        }
    }
    traceoptions {
        file <filename> <files number> <match regular-expression> <size bytes> <world-readable | no-
world-readable>;
        no-remote-trace;
```

```
    }
}
```

NOTE: When a route pointing to more than one services PIC is available, and with application layer gateways (ALGs) configured, you must always configure the distribution of traffic across PICs based on the source IP address by including the `family inet layer-3-services source-address` statement at the [`edit forwarding-options enhanced-hash-key services-loadbalancing`] hierarchy level for IPv4 traffic and the `family inet6 layer-3-services source-address` statement at the [edit forwarding-options enhanced-hash-key services-loadbalancing] hierarchy level for IPv6 traffic. With ALGs used to manage a parent-child relationship of sessions, both the parent and the child sessions must be processed by the same type of services PIC.

Configuring IPv4 and IPv6 Accounting

IPv4 and IPv6 accounting are disabled by default. But you can enable the accounting by including the `route-accounting` statements at the [`edit forwarding-options family inet4`] and [`edit forwarding-options family inet6`] hierarchy level, as shown here:

```
[edit]
forwarding-options {
    family inet4 {
        route-accounting;
    }
]
```

```
[edit]
forwarding-options {
    family inet6 {
        route-accounting;
    }
]
```

To view the IPv4 or IPv6 statistics for a given physical or logical interface, run the operational command `show interfaces extensive` *`interface`* `| find IPv4` or `show interfaces extensive` *`interface`* `| find IPv6`. Note that the

output displays packet and byte counts for transit traffic only. Locally generated packets are not included in the metrics.

```
show interfaces ge-2/0/9 detail | find IPv6
IPv6 transit statistics:
Input  bytes  :          8576802312
Output bytes  :        8991637500
Input  packets:        5787313
Output packets:      5994425
```

Follow the guidelines given below to configure IPv4 and IPv6 statistics accounting:

- The transit rate is calculated in software and not in hardware. So, you can expect some deviation from the line rate. You do not receive the expected rate with lower fps (for example, 5 fps).
- The traffic must be IPv4 or IPv6 at both ingress and egress. Any form of encapsulated traffic or traffic translating to IPv4 is not supported on PTX routers.
- If the GRE tunnel is IPv4 or IPv6 based, this traffic is not accounted in ingress statistics.
- Statistics accounting is supported only for `et` and `ae` interfaces on PTX routers.
- The virtual, software, pseudo, Layer 2, and special type Layer 3 virtual interfaces are not supported. For example, interfaces such as the following:
  - `irb`
  - `rvi`
  - `lsi`
  - `fti`
  - `vtep`
  - `vt`
- IPv4 transit statistics is not displayed on Routine Engine show commands.
- No interoperability between SCU and IPv4 statistics accounting on PTX routers.

  The following functionalities are not supported:
  - lo0 traffic
  - IPv4 statistics query over SNMP
  - Aggregated Ethernet child link IPv4 statistics accounting

RELATED DOCUMENTATION

route-accounting | 300

show interfaces statistics

Configuring Discard Accounting

On routing platforms containing a Monitoring Services PIC or an Adaptive Services PIC, you can configure accounting for traffic passing through the routing platform.

To configure discard accounting, include the `accounting group` *`group-name`* statement at the `[edit forwarding-options]` hierarchy level:

```
[edit forwarding-options]
accounting group-name {
    output {
        cflowd [ hostnames ] {
            aggregation {
                autonomous-system;
                destination-prefix;
                protocol-port;
                source-destination-prefix {
                    caida-compliant;
                }
                source-prefix;
            }
            autonomous-system-type (origin | peer);
            port port-number;
            version format;
        }
        flow-active-timeout seconds;
        flow-inactive-timeout seconds;
        interface {
            engine-id number;
            engine-type number;
            source-address address;
        }
    }
}
```

To configure the output flow aggregation, include the `cflowd` statement. For more information about flow aggregation, see "Directing Traffic Sampling Output to a Server Running the cflowd Application" on page 12. To configure the interval before exporting an active flow, include the `flow-active-timeout` statement. The default value for **flow-active-timeout** is **1800** seconds. To configure the interval before a flow is considered inactive, include the `flow-inactive-timeout` statement. The default value for **flow-inactive-timeout** is 60 seconds. To configure the interface that sends out monitored information, include the `interface` statement. Discard accounting is supported for the Monitoring Services PIC only.

When you apply a firewall filter to a loopback interface, the filter might block responses from the Monitoring Services PIC. To allow responses from the Monitoring Services PIC to pass through for accounting purposes, configure a term in the firewall filter to include the Monitoring Services PIC IP address. For more detailed information about configuring firewall filters, see *Guidelines for Configuring Firewall Filters* and *Guidelines for Applying Standard Firewall Filters*.

You can use discard accounting for passive and active flow monitoring.

RELATED DOCUMENTATION

- Monitoring, Sampling, and Collection Services Interfaces User Guide
- Junos OS Class of Service User Guide for Routing Devices

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

You can use flow monitoring to help with network administration. Active flow monitoring on PTX Series routers allows you to collect sampled packets, then the router does GRE encapsulation of the packets and sends them to a remote server for flow processing. The GRE encapsulation includes an interface index and GRE key field. The GRE encapsulation removes MPLS tags. You configure one or more port-mirroring instances to define which traffic to sample and configure a server to receive the GRE encapsulated packets. You configure a firewall filter on interfaces where you want to capture flows. You can configure as many as 48 port-mirroring instances.

To configure the router to do GRE encapsulation of sampled packets and send them to a remote server for flow processing:

1. Configure one or more server profiles that specify a host where GRE encapsulated sampled packets are sent, and optionally, a source address to include in the header of each sampled packet.

a. Specify a name for each server profile and an IP address of the host where sampled packets are sent:

```
[edit services hosted-services]
user@host# set server-profile server-profile-name server-address ipv4-address
```

b. (Optional) For each server profile, specify a source address to include in the header of each sampled packet:

```
[edit services hosted-services server-profile server-profile-name]
user@host# set client-address ipv4-address
```

NOTE: The default client address is 0.0.0.0. You must specify an IPv4 address as the client address. You can also specify the loopback address or management interface address as the client address.

2. Configure one or more port-mirroring instances.

a. Specify a name for each port-mirroring instance:

```
[edit forwarding-options port-mirroring]
user@host# set instance instance-name
```

NOTE: You can configure a maximum of 48 port-mirroring instances.

b. Specify a protocol family for each port-mirroring instance:

```
[edit forwarding-options port-mirroring instance instance-name]
user@host# set family (inet | inet6 )
```

3. To set the ratio of the number of packets to sample, specify a value from 1 through 65,535 for each port-mirroring instance:

```
[edit forwarding-options port-mirroring instance instance-name input]
user@host# set rate number
```

NOTE: You must specify a value for the `rate` statement. The default value is zero, which effectively disables sampling. If, for example, you specify a rate value of 4, every fourth packet (1 packet out of 4) is sampled.

4. (Optional) Specify the number of samples to collect after the initial trigger event for each port-mirroring instance:

```
[edit forwarding-options port-mirroring instance instance-name input]
user@host# set run-length number
```

NOTE: The default value is zero. You can specify a number up to 20.

5. To designate a host where sampled traffic is sent, specify the name of server profile configured at the `[edit services hosted-services]` hierarchy level for each port-mirroring instance:

```
[edit forwarding-options port-mirroring instance instance-name family ( inet | inet6) output]
user@host# set server-profile server-profile-name
```

6. Configure one or more firewall filters.
   a. For each firewall filter, specify a protocol family, filter name, and match conditions:

```
[edit firewall]
user@host# set filter family (inet | inet6) filter filter-name term term-name from match-condtions
```

   b. For each firewall filter you configure, specify the name of a port-mirroring instance you configured at the `[edit forwarding-options]` hierarchy level as a nonterminating action so that the traffic that matches that instance is sampled:

```
[edit firewall family (inet | inet6) filter filter-name term term-name]
user@host# set then port-mirroring instance instance-name
```

7. Apply each firewall filter to an interface to evaluate incoming traffic:

```
[edit interfaces interface-name unit logical-unit-number]
user@host# set family (inet | inet6) filter input firewall-filter-name
```

> **NOTE**: Active flow monitoring is supported only on incoming traffic. You cannot apply firewall filters to evaluate outgoing traffic.

8. Configure the remote server, where GRE encapsulated packets are sent, to perform flow processing.

RELATED DOCUMENTATION

Configuring Port Mirroring | 37

hosted-services | 218

port-mirroring | 291

server-profile (Active Flow Monitoring) | 315

Firewall Filter Nonterminating Actions

Configuring Passive Flow Monitoring

On routing platforms containing the Monitoring Services PIC or the Monitoring Services II PIC, you can configure flow monitoring for traffic passing through the routing platform. This type of monitoring method is passive monitoring.

To configure flow monitoring, include the `monitoring` statement at the `[edit forwarding-options` hierarchy level:

```
[edit forwarding-options]
monitoring group-name {
    family inet {
        output {
            cflowd hostname {
                port port-number;
            }
            export-format cflowd-version-5;
```

```
            flow-active-timeout seconds;
            flow-export-destination {
                cflowd-collector;
            }
            flow-inactive-timeout seconds;
            interface interface-name {
                engine-id number;
                engine-type number;
                input-interface-index number;
                output-interface-index number;
                source-address address;
            }
        }
    }
}
```

To configure a passive monitoring group, include the `monitoring` statement and specify a group name. To configure monitoring on a specified address family, include the `family` statement and specify an address family. To specify an interface to monitor incoming traffic, include the `input` statement. To configure the monitoring information that is sent out, include the `output` statement. To configure the output flow aggregation, include the `cflowd` statement. For more information about flow aggregation, see "Directing Traffic Sampling Output to a Server Running the cflowd Application" on page 12. To specify the format of the monitoring information sent out, include the `export-format` statement and specify a version number. To configure the interval before exporting an active flow, include the `flow-active-timeout` statement. The default value for **flow-active-timeout** is 1800 seconds. To enable flow collection, include the `flow-export-destination` statement. To configure the interval before a flow is considered inactive, include the `flow-inactive-timeout` statement. The default value for **flow-inactive-timeout** is 60 seconds. To configure the interface that sends out the monitored information, include the `interface` statement. Flow monitoring is supported for Monitoring Services PIC interfaces only.

When you apply a firewall filter to a loopback interface, the filter might block responses from the Monitoring Services PIC. To allow responses from the Monitoring Services PIC to pass through for monitoring purposes, configure a term in the firewall filter to include the Monitoring Services PIC's IP address. For more detailed information about configuring firewall filters, see *Guidelines for Configuring Firewall Filters* and *Guidelines for Applying Standard Firewall Filters*.

RELATED DOCUMENTATION

- Monitoring, Sampling, and Collection Services Interfaces User Guide
- Junos OS Class of Service User Guide for Routing Devices

Configuring Port Mirroring

IN THIS SECTION

- Port Mirroring Configuration Guidelines | 38
- Configuring Port Mirroring | 39
- Configuring Multiple Port-Mirroring Instances | 40

Port mirroring is the ability of a router to send a copy of an IPv4 or IPv6 packet to an external host address or a packet analyzer for analysis. Port mirroring is different from traffic sampling. In traffic sampling, a sampling key based on the packet header is sent to the Routing Engine. There, the key can be placed in a file, or cflowd packets based on the key can be sent to a cflowd server. In port mirroring, the entire packet is copied and sent out through a next-hop interface.

One application for port mirroring sends a duplicate packet to a virtual tunnel. A next-hop group can then be configured to forward copies of this duplicate packet to several interfaces. For more information about next-hop groups, see "Configuring Next-Hop Groups to Use Multiple Interfaces to Forward Packets Used in Port Mirroring" on page 42.

All M Series Multiservice Edge Routers, T Series Core Routers, and MX Series 5G Universal Routing Platforms support port mirroring for IPv4 or IPv6. The M120, M320, and MX Series routers support port mirroring for IPv4 and IPv6 simultaneously.

Port mirroring for VPLS traffic is supported on M7i and M10i routers configured with an Enhanced CFEB (CFEB-E), on M120 routers, on M320 routers configured with an Enhanced III Flexible PIC Concentrators (FPCs), and MX Series routers.

In Junos OS Release 9.3 and later, port mirroring is supported for Layer 2 traffic on MX Series routers. For information about how to configure port mirroring for Layer 2 traffic, see the Junos OS Layer 2 Switching and Bridging Library for Routing Devices.

In Junos OS Release 9.6 and later, port mirroring is supported for Layer 2 VPN traffic on M120 routers and M320 routers configured with an Enhanced III FPC. You can also set the maximum length of the mirrored packet. When set, the mirrored packet is truncated to the specified length.

In the MPCs on M Series and MX Series routers, GRE and MPLS header information is not contained in the port-mirrored traffic corresponding to MPLS packets transmitted through IP-GRE tunnels.

Port Mirroring Configuration Guidelines

When configuring port mirroring, the following restrictions apply:

- Only transit data is supported.
- You can configure either IPv4 or IPv6 port mirroring but not both on M Series routers, except for the M120 and M320 routers, which support port mirroring for IPv4 and IPv6 simultaneously.
- You can configure port mirroring for IPv4 and IPv6 simultaneously on the M120 and M320 routers and the MX Series routers.
- Port mirroring in the ingress and egress direction is not supported for link services IQ (lsq-) interfaces.
- Ingress filtering of multicast packets is supported on all Dense Port Concentrators (DPCs) in MX Series routers. Egress filtering of multicast packets is supported for interfaces on MPCs in MX Series routers only. Filtering of multicast packets based on destination address is not supported on M Series routers or T Series routers and is not supported for interfaces on I-chip ASIC-based DPCs in MX Series routers.

  For Layer 3 port mirroring (`family inet` and `family inet6`), if the traffic being mirrored is multicast (in other words, if the packet's destination IP address is a multicast address), the destination MAC address in the mirrored copy corresponds to this multicast destination IP address and not to the unicast MAC address specified in the `[edit forwarding-options port-mirroring family (inet | inet6) output]` configuration.
- By default, firewall filters cannot be applied to port-mirroring destination interfaces. To enable port-mirroring destination interfaces to support firewall filters, use the `no-filter-check` statement to disable filter checking on the interfaces. You can include the `no-filter-check` statement at the following hierarchy levels:
  - `[edit forwarding-options port-mirroring family (inet | inet6 | ccc | vpls) output]`
  - `[edit forwarding-options port-mirroring instance` *`instance-name`* `family (inet | ccc | vpls) output]`
- You must include a firewall filter with both the **accept** action and the **port-mirror** action modifier on the inbound interface.
- The interface you configure for port mirroring should not participate in any kind of routing activity.
- The destination address you specify should not have a route to the ultimate traffic destination. For example, if the sampled IPv4 packets have a destination address of **192.68.9.10** and the port-mirrored traffic is sent to **192.68.20.15** for analysis, the device associated with the latter address should not know a route to **192.68.9.10**. Also, it should not send the sampled packets back to the source address.

- On all routers except the MX Series router, you can configure only one port-mirroring interface per router. If you include more than one interface in the `port-mirroring` statement, the previous one is overwritten. MX Series routers support more than one port-mirroring interface per router.
- You can configure multiple port-mirroring instances on the M120, M320, and MX Series routers.
- You can specify both host (cflowd) sampling and port mirroring in the same configuration. You can perform RE-sampling and port mirroring actions simultaneously. However, you cannot perform PIC-sampling and port mirroring actions simultaneously.
- In typical applications, you send the sampled packets to an analyzer or a workstation for analysis, not to another router. If you must send this traffic over a network, you should use tunnels.

Configuring Port Mirroring

To configure port mirroring, include the `port-mirroring` statement at the `[edit forwarding-options]` hierarchy level:

```
[edit forwarding-options]
port-mirroring {
    family (ccc | inet | inet6 | vpls) {
        output {
            interface interface-name {
                next-hop address;
            }
            no-filter-check;
        }
        input {
            maximum-packet-length bytes;
            rate number;
            run-length number;
        }
    }
}
```

Configuring the Port-Mirroring Address Family and Interface

To configure port mirroring, include the `port-mirroring` statement. To configure the address family type of traffic to sample, include the `family` statement. To configure the rate of sampling, length of sampling, and the maximum size for the mirrored packet, include the `input` statement. To specify on which interface to

send duplicate packets and the next-hop address to send packets, include the `output` statement. To determine whether there are any filters on the specified interface, include the `no-filter-check` statement.

For information about the **rate** and `run-length` statements, see "Configuring Traffic Sampling" on page 6 .

Configuring Multiple Port-Mirroring Instances

In Junos OS Release 9.5 and later, you can configure multiple port-mirroring instances on the M120, M320, and MX Series routers. On the M120 router, you can associate each instance with a specific Forwarding Engine Board (FEB). You cannot associate a port-mirroring instance with an FEB configured as a backup FEB. On the M320 router, you can associate each instance with a specific Flexible PIC Concentrator (FPC). Associating a port-mirroring instance with an FPC or an FEB enables you to mirror packets to different destinations. Multiple port-mirroring instances are also supported on MX Series routers. For information about configuring multiple port-mirroring instances on MX Series routers, see the Junos OS Layer 2 Switching and Bridging Library for Routing Devices.

NOTE: In MX80 and MX104 routers, port-mirroring instances should always be associated with FPC 0, because associating port-mirroring instances to FPC 1 or FPC 2 can result in inconsistent behavior due to the underlying architecture.

To configure a port-mirroring instance, include the `instance` *`port-mirroring-instance`* statement at the `[edit forwarding-options port-mirroring]` hierarchy level:

```
[edit forwarding-options port-mirroring]
instance port-mirroring-instance-name {
    family (ccc | inet | inet6 | vpls) {
        output {
            interface interface-name {
                next-hop address;
            }
            no-filter-check;
        }
    }
    input {
        maximum-packet-length bytes;
        rate number;
        run-length number;
```

```
    }
}
```

Configuring Port-Mirroring Instances

You can configure multiple port-mirroring instances. Specify a unique ***port-mirroring-instance-name*** for each instance you configure.

Associating a Port-Mirroring Instance on M320 Routers

You can associate a port-mirroring instance with a specific FPC on an M320 router or with a specific FEB on an M120 router. You can associate only one port-mirroring instance with each FPC on an M320 router or with each FEB on an M120 router. On an M120 router, you cannot associate a port-mirroring instance with a FEB configured as a backup FEB.

To associate a port-mirroring instance with an FPC on an M320 router, include the `port-mirror-instance` *`port-mirroring-instance-name`* statement at the `[edit chassis fpc` *`slot-number`*`]` hierarchy level:

```
[edit chassis]
fpc slot-number {
    port-mirror-instance port-mirroring-instance-name;
}
```

For ***slot-number***, specify the slot number of the FPC you want to associate with the port-mirroring instance. For ***port-mirroring-instance-name***, specify the name of a port-mirroring instance you configured at the `[edit forwarding-options port-mirroring]` hierarchy level. For more information about configuring an FPC on an M320 router, see the Junos OS Administration Library for Routing Devices.

Associating a Port-Mirroring Instance on M120 Routers

To associate a port-mirroring instance with a FEB on an M120 router, include the `port-mirror-instance` *`port-mirroring-instance-name`* statement at the `[edit chassis feb` *`slot-number`*`]` hierarchy level:

```
[edit chassis]
feb slot-number {
    port-mirror-instance port-mirroring-instance-name;
}
```

For ***slot-number***, specify the slot number of the FEB you want to associate with the port-mirroring instance. For ***port-mirroring-instance-name***, specify the name of a port-mirroring instance you

configured at the [edit forwarding-options port-mirroring] hierarchy level. For information about configuring FEB redundancy on an M120 router, see the Junos OS High Availability User Guide. For information about configuring FPC-to-FEB connectivity on an M120 router, see the Junos OS Administration Library for Routing Devices.

Configuring MX Series 5G Universal Routing Platforms and M120 Routers to Mirror Traffic Only Once

On MX Series and M120 routers only, you can configure port mirroring so that the router mirrors traffic only once. If you configure port mirroring on both ingress and egress interfaces, the same packet could be mirrored twice. To mirror packets only once and prevent the router from sending duplicate sampled packets to the same mirroring destination, include the mirror-once statement at the [edit forwarding-options port-mirroring] hierarchy level:

```
[edit forwarding-options port-mirroring]
mirror-once;
```

NOTE: The mirror-once statement is supported only in the global port-mirroring instance.

Configuring Next-Hop Groups to Use Multiple Interfaces to Forward Packets Used in Port Mirroring

Next-hop groups allow you to include multiple interfaces used to forward duplicate packets used in port mirroring.

To configure a next-hop group, include the next-hop-group statement at the [edit forwarding-options] hierarchy level:

```
[edit forwarding-options]
next-hop-group [ group-names ] {
    interface interface-name {
        next-hop [ addresses ];
```

```
    }
}
```

or

```
[edit forwarding-options port-mirroring family inet6 output]
next-hop-group group-name{
    group-type inet6;
    interface interface-name {
        next-hop ipv6-address;
    }
    next-hop-subgroup group-name{
        interface interface-name {
            next-hop ipv6-address;
        }
    }
}
```

You can specify one or more group names. To configure the interface that sends out sampled information, include the `interface` statement and specify an interface. To specify a next-hop address to send sampled information, include the `next-hop` statement and specify an IP address.

Next-hop groups have the following restrictions:

- Starting with release 14.2, next-hop groups are supported for M Series and MX Series routers only.
- Next-hop groups support up to 16 next-hop addresses.
- You can configure up to 30 next-hop groups.
- Each next-hop group must have at least two next-hop addresses.
- When a firewall filter with next-hop-group action is applied on an interface in egress, the redirected copy does not retain any packet headers added while forwarding the packet to that interface. For example, if a filter with action next-hop-group is applied in egress of a GRE interface, the redirected copies received on the next-hop-group member interfaces do not contain a GRE header.

Next-hop groups can be used for port mirroring.

Release History Table

Release	Description
14.2	Starting with release 14.2, next-hop groups are supported for M Series and MX Series routers only.

RELATED DOCUMENTATION

Configuring Port Mirroring | 37

Defining a Port-Mirroring Firewall Filter

Starting with release 14.2, on routers containing an Internet Processor II application-specific integrated circuit (ASIC) or T Series Internet Processor, you can send a copy of an IP version 4 (IPv4) or IP version 6 (IPv6) packet from the router to an external host address or a packet analyzer for analysis. This is known as *port mirroring*.

Port mirroring is different from traffic sampling. In traffic sampling, a sampling key based on the IPv4 header is sent to the Routing Engine. There, the key can be placed in a file, or cflowd packets based on the key can be sent to a cflowd server. In port mirroring, the entire packet is copied and sent out through a next-hop interface.

You can configure simultaneous use of sampling and port mirroring, and set an independent sampling rate and run-length for port-mirrored packets. However, if a packet is selected for both sampling and port mirroring, only one action can be performed and port mirroring takes precedence. For example, if you configure an interface to sample every packet input to the interface and a filter also selects the packet to be port mirrored to another interface, only the port mirroring would take effect. All other packets not matching the explicit filter port-mirroring criteria continue to be sampled when forwarded to their final destination.

Firewall filters provide a means of protecting your router from excessive traffic transiting the router to a network destination or destined for the Routing Engine. Firewall filters that control local packets can also protect your router from external incidents.

You can configure a *firewall filter* to do the following:

- Restrict traffic destined for the Routing Engine based on its source, protocol, and application.
- Limit the traffic rate of packets destined for the Routing Engine to protect against flood, or denial-of-service (DoS) attacks.
- Address special circumstances associated with fragmented packets destined for the Routing Engine. Because the device evaluates every packet against a firewall filter (including fragments), you must configure the filter to accommodate fragments that do not contain packet header information. Otherwise, the filter discards all but the first fragment of a fragmented packet.

For information about configuring firewall filters in general (including in a Layer 3 environment), see *Stateless Firewall Filter Overview* and *How Standard Firewall Filters Evaluate Packets* in the Routing Policies, Firewall Filters, and Traffic Policers User Guide.

To define a firewall filter with a port-mirroring action:

1. Prepare traffic for port mirroring by including the `filter` statement at the `[edit firewall family (inet | inet6)]` hierarchy level.

```
filter filter-name;
```

   This filter at the `[edit firewall family (inet | inet6)]` hierarchy level selects traffic to be port-mirrored:

```
filter filter-name {
    term term-name {
        then {
            port-mirror;
            accept;
        }
    }
}
```

2. Enable configuration of firewall filters.

```
[edit]
user@host# edit firewall family family
```

   The value of the *family* option can be `inet` or `inet6`.

3. Enable configuration of a firewall filter ***filter-name***.

```
[edit firewall family family]
user@host# edit filter filter-name
```

4. Enable configuration of a firewall filter term ***filter-term-name***.

```
[edit firewall family family filter filter-name]
user@host# edit term filter-term-name
```

   For more information about firewall filter terms, see *Guidelines for Configuring Firewall Filters* in the Junos OS Routing Policies, Firewall Filters and Traffic Policers User Guide for Routing Devices.

5. Specify the firewall filter match conditions based on the route source address to mirror a subset of the sampled packets.

For information about configuring firewall filter match conditions, see *Firewall Filter Match Conditions Based on Numbers or Text Aliases*, *Firewall Filter Match Conditions Based on Bit-Field Values*, *Firewall Filter Match Conditions Based on Address Fields*, and *Firewall Filter Match Conditions Based on Address Classes*in the Junos OS Routing Policies, Firewall Filters and Traffic Policers User Guide for Routing Devices.

6. Enable configuration of the ***action*** and ***action-modifier*** to apply to the matching packets.

```
[edit firewall family family filter filter-name term filter-term-name]
user@host# edit then
```

7. Specify the actions to be taken on matching packets.

```
[edit firewall family family filter filter-name term filter-term-name then]
user@host# set action
```

The recommended value for the ***action*** is **accept**. If you do not specify an action, or if you omit the `then` statement entirely, all packets that match the conditions in the `from` statement are accepted.

8. Specify port-mirror as the ***action-modifier***.

When the filter action is `port-mirror`, the packet is copied to a local interface for local or remote monitoring.

```
[edit firewall family family filter filter-name term filter-term-name then]
user@host# set port-mirror
```

9. Verify the minimum configuration of the firewall filter.

```
[edit firewall ... ]
user@host# top
[edit]
user@host# show firewall

family (inet | inet6) { # Type of packets to mirror
    filter filter-name { # Firewall filter name
        term filter-term-name {
            from { # Do not specify match conditions based on route source address
            }
            then {
                port-mirror;
                accept;
            }
```

```
        }
    }
}
```

Release History Table

Release	Description
14.2	Starting with release 14.2, on routers containing an Internet Processor II application-specific integrated circuit (ASIC) or T Series Internet Processor, you can send a copy of an IP version 4 (IPv4) or IP version 6 (IPv6) packet from the router to an external host address or a packet analyzer for analysis.

RELATED DOCUMENTATION

Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers

Example: Configuring Multiple Port Mirroring with Next-Hop Groups on M, MX and T Series Routers

Defining a Next-Hop Group on MX Series Routers for Port Mirroring

Starting with release 14.2, on routers containing an Internet Processor II application-specific integrated circuit (ASIC) or T Series Internet Processor, you can send a copy of an IP version 4 (IPv4) or IP version 6 (IPv6) packet from the router to an external host address or a packet analyzer for analysis. This is known as *port mirroring*.

Port mirroring is different from traffic sampling. In traffic sampling, a sampling key based on the IPv4 header is sent to the Routing Engine. There, the key can be placed in a file, or cflowd packets based on the key can be sent to a cflowd server. In port mirroring, the entire packet is copied and sent out through a next-hop interface.

You can configure simultaneous use of sampling and port mirroring, and set an independent sampling rate and run-length for port-mirrored packets. However, if a packet is selected for both sampling and port mirroring, only one action can be performed, and port mirroring takes precedence. For example, if you configure an interface to sample every packet input to the interface and a filter also selects the packet to be port mirrored to another interface, only the port mirroring takes effect. All other packets not matching the explicit filter port-mirroring criteria continue to be sampled when forwarded to their final destination.

Next-hop groups allow you to include port mirroring on multiple interfaces.

On MX Series routers, you can mirror tunnel interface input traffic to multiple destinations. To this form of multipacket port mirroring, you specify two or more destinations in a next-hop group, define a firewall filter that references the next-hop group as the filter action, and then apply the filter to a logical tunnel interface `lt-`) or virtual tunnel interfaces (`vt-` on the MX Series router.

To define a next-hop group for a Layer 2 port-mirroring firewall filter action:

1. Enable the configuration of forwarding options.

```
[edit]
user@host set forwarding-options port-mirroring family (inet | inet6) output
```

2. Enable configuration of a next-hop-group for Layer 2 port mirroring.

```
[edit forwarding-options port-mirroring ... family (inet | inet6) output]
user@host# set next-hop-group next-hop-group-name
```

3. Specify the type of addresses to be used in the next-hop group configuration.

```
[edit forwarding-options port-mirroring ... family (inet | inet6) output next-hop-group next-hop-group-name]
user@host# set group-type inet6
```

4. Specify the interfaces of the next-hop route.

```
[edit forwarding-options port-mirroring ... family (inet | inet6) output next-hop-group next-hop-group-name]
user@host# set interface logical-interface-name-1
user@host# set interface logical-interface-name-2
```

or

```
[edit forwarding-options port-mirroring ... family (inet | inet6) output next-hop-group next-hop-group-name]
user@host# set interface interface-name next-hop next-hop-address
```

The MX Series router supports up to 30 next-hop groups. Each next-hop group supports up to 16 next-hop addresses. Each next-hop group must specify at least two addresses. The *next-hop-address* can be an IPv4 or IPv6 address.

5. (Optional) Specify the next-hop subgroup.

```
[edit forwarding-options port-mirroring ... family (inet | inet6) output next-hop-group next-hop-group-name]
user@host# set next-hop-subgroup subgroup-name interface interface-name next-hop next-hop-address
```

6. Verify the configuration of the next-hop group.

```
[edit forwarding-options port-mirroring ... family (inet | inet6) output next-hop-group next-hop-group-name]
user@host# top
[edit]
user@host# show forwarding-options

...
next-hop-group next-hop-group-name {
    group-type inet6;
    interface logical-interface-name-1;
    interface interface-name{
        next-hop next-hop-address;
    }
    next-hop-subgroup subgroup-name{
        interface interface-name{
            next-hop next-hop-address;
        }
    }
}
...
```

Release History Table

Release	Description
14.2	Starting with release 14.2, on routers containing an Internet Processor II application-specific integrated circuit (ASIC) or T Series Internet Processor, you can send a copy of an IP version 4 (IPv4) or IP version 6 (IPv6) packet from the router to an external host address or a packet analyzer for analysis.

CHAPTER 4

Configuring Forwarding Table Filters to Efficiently Route Traffic

Configuring Forwarding Table Filters | 51

Forwarding Table Filters for Routing Instances on ACX Series Routers | 53

Applying Forwarding Table Filters | 54

Configuring Forwarding Table Filters

Forwarding table filters are defined the same as other firewall filters, but you apply them differently:

- Instead of applying forwarding table filters to interfaces, you apply them to forwarding tables, each of which is associated with a routing instance and a virtual private network (VPN).
- Instead of applying input and output filters by default, you can apply an input forwarding table filter only.

All packets are subjected to the input forwarding table filter that applies to the forwarding table. A forwarding table filter controls which packets the router accepts and then performs a lookup for the forwarding table, thereby controlling which packets the router forwards on the interfaces.

When the router receives a packet, it determines the best route to the ultimate destination by looking in a forwarding table, which is associated with the VPN on which the packet is to be sent. The router then forwards the packet toward its destination through the appropriate interface.

NOTE: For transit packets exiting the router through the tunnel, forwarding table filtering is not supported on the interfaces you configure as the output interface for tunnel traffic.

A forwarding table filter allows you to filter data packets based on their components and to perform an action on packets that match the filter; it essentially controls which bearer packets the router accepts and forwards. To configure a forwarding table filter, include the `firewall` statement at the `[edit]` hierarchy level:

```
[edit]
firewall {
    family family-name {
        filter filter-name {
            term term-name {
                from {
                    match-conditions;
                }
                then {
                    action;
                    action-modifiers;
                }
            }
        }
```

```
    }
}
```

family-name is the family address type: IPv4 (**inet**), IPv6 (**inet6**), Layer 2 traffic (**bridge**), or MPLS (**mpls**).

term-name is a named structure in which match conditions and actions are defined.

match-conditions are the criteria against which a bearer packet is compared; for example, the IP address of a source device or a destination device. You can specify multiple criteria in a match condition.

action specifies what happens if a packet matches all criteria; for example, the gateway GPRS support node (GGSN) accepting the bearer packet, performing a lookup in the forwarding table, and forwarding the packet to its destination; discarding the packet; and discarding the packet and returning a rejection message.

action-modifiers are actions that are taken in addition to the GGSN accepting or discarding a packet when all criteria match; for example, counting the packets and logging a packet.

To create a forwarding table, include the `instance-type` statement with the **forwarding** option at the `[edit routing-instances` *`instance-name`*`]` hierarchy level:

```
[edit]
routing-instances instance-name {
    instance-type forwarding;
}
```

To apply a forwarding table filter to a VPN routing and forwarding (VRF) table, include the **filter** and `input` statements at the `[edit routing-instance` *`instance-name`* `forwarding-options family` *`family-name`*`]` hierarchy level:

```
[edit routing-instances instance-name]
instance-type forwarding;
forwarding-options {
    family family-name {
        filter {
            input filter-name;
        }
    }
}
```

To apply a forwarding table filter to a forwarding table, include the **filter** and `input` statements at the `[edit forwarding-options family family-name]` hierarchy level:

```
[edit forwarding-options family family-name]
filter {
    input filter-name;
}
```

To apply a forwarding table filter to the default forwarding table **inet.0**, which is not associated with a specific routing instance, include the **filter** and `input` statements at the `[edit forwarding-options family inet]` hierarchy level:

```
[edit forwarding-options family inet]
filter {
    input filter-name;
}
```

RELATED DOCUMENTATION

- *Guidelines for Configuring Firewall Filters*
- *Guidelines for Applying Standard Firewall Filters*
- Applying Forwarding Table Filters | 54

Forwarding Table Filters for Routing Instances on ACX Series Routers

Forwarding table filter is a mechanism by which all the packets forwarded by a certain forwarding table are subjected to filtering and if a packet matches the filter condition, the configured action is applied on the packet. You can use the forwarding table filter mechanism to apply a filter on all interfaces associated with a single routing instance with a simple configuration. You can apply a forwarding table filter to a routing instance of type forwarding and also to the default routing instance `inet.0`. To configure a forwarding table filter, include the `filter filter-name` statement at the `[edit firewall family <inet | inet6>]` hierarchy level.

The following limitations apply to forwarding table filters configured on routing instances:

- You cannot attach the same filter to more than one routing instance.
- You cannot attach the same filter at both the [edit interfaces *interface-name* family <inet | inet6> filter input *filter-name*] and [edit routing-instances *instance-name* forwarding-options family <inet | inet6> filter input *filter-name*] hierarchy level.
- You cannot attach a filter that is either interface-specific or a physical interface filter.
- You cannot attach a filter to the egress direction of routing instances.

RELATED DOCUMENTATION

> *Configuring Forwarding Table Filters*

Applying Forwarding Table Filters

A forwarding table filter allows you to filter data packets based on their components and perform an action on packets that match the filter. You can apply a filter on the ingress or egress packets of a forwarding table. You configure the filter at the [edit firewall family *family-name*] hierarchy level; for more information, see *Configuring Forwarding Table Filters*.

To apply a forwarding table filter on ingress packets of a forwarding table, include the **filter** and input statements at the [edit forwarding-options family *family-name*] hierarchy level:

```
[edit forwarding-options family family-name]
filter {
    input filter-name;
}
```

You can filter based upon destination-class information by applying a firewall filter on the egress packets of the forwarding table. By applying firewall filters to packets that have been forwarded by a routing table, you can match based on certain parameters that are decided by the route lookup. For example, routes can be classified into specific destination and source classes. Firewall filters used for policing and mirroring are able to match based upon these classes.

To apply a firewall filter on egress packets of a forwarding table, include the **filter** and `output` statements at the `[edit forwarding-options family family-name]` hierarchy level:

```
[edit forwarding-options family family-name]
filter {
    output filter-name;
}
```

NOTE: You cannot have a firewall filter that includes an `interface-group` match condition if you are also using an egress forwarding table filter. This is because the `interface-group` match condition uses the logical interface on which the packet was received to match the interface group (or set of interface groups), while the forwarding table filters apply only to local host traffic and transit packets.

To apply a forwarding table filter to a flood table, include the **flood** and `input` statements at the `[edit forwarding-options family family-name]` hierarchy level as shown below. The `flood` statement is valid for the **vpls** protocol family only.

```
[edit forwarding-options family vpls]
flood {
    input filter-name;
}
```

On the MX Series router only, to apply a forwarding table filter for a virtual switch, include the **filter** and `input` statements at the `[edit routing-instances routing-instance-name bridge-domains bridge-domain-name forwarding-options]` hierarchy level:

```
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name forwarding-
options]
    filter {
        input filter-name;
    }
```

For more information about how to configure a virtual switch, see the Junos OS Layer 2 Switching and Bridging Library for Routing Devices.

On MX Series 3D Universal Edge Routers, you can apply a forwarding table filter by using the `soure-checking` statement at the `[edit forwarding-options family inet6]` hierarchy level:

```
[edit forwarding-options family inet6]
    family inet6 {
        source-checking;
    }
}
```

This discards IPv6 packets when the source address type is unspecified, loopback, multicast or link-local.

RFC 4291, *IP Version 6 Addressing Architecture*, refers to four address types that require special treatment when they are used as source addresses. The four address types are:

- Unspecified
- Loopack
- Multicast
- Link-Local Unicast

The loopback and multicast addresses must never be used as a source address in IPv6 packets. The unspecified and link-local addresses can be used as source addresses but routers must never forward packets that have these addresses as source addresses. Typically, packets that contain unspecified or link-local addresses as source addresses are delivered to the local host. If the destination is not the local host, then the packet must not be forwarded. Configuring this statement filters or discards IPv6 packets of these four address types.

NOTE: For T Series routers other than T4000, a packet forwarded by the forwarding table reaches the egress forwarding table filter irrespective of whether the packet is actually forwarded by the forwarding table or not. The packet reaches the egress filter even if the route points to reject or discard next hops.

On T4000 Type 5 Flexible PIC Concentrator (FPC), the packet reaches the egress filter only if it is forwarded by the forwarding table.

NOTE: The egress forwarding table filter is applied on the ingress interface of the FPC. If different packets to the same destination arrive on different FPCs, they might encounter different policers.

NOTE: In versions 14.2 and prior, the egress forwarding table filter is not supported for the J Series Service Routers.

NOTE: In Junos OS Release 8.4 and later, you can no longer configure this output statement for VPLS. You can continue to configure ingress forwarding table filters with the input statement at the `[edit forwarding-options family vpls filter]` hierarchy level.

Release History Table

Release	Description
14.2	In versions 14.2 and prior, the egress forwarding table filter is not supported for the J Series Service Routers.
8.4	In Junos OS Release 8.4 and later, you can no longer configure this output statement for VPLS.

5
CHAPTER

Configuring Forwarding Options for Load Balancing Traffic

Configuring Load Balancing for Ethernet Pseudowires | 59

Configuring Load-Balance Groups | 60

Understanding the Algorithm Used to Load Balance Traffic on MX Series Routers | 61

Understanding Per-Packet Load Balancing | 75

Configuring Per-Packet Load Balancing | 77

Understanding Load Balancing for BGP Traffic with Unequal Bandwidth Allocated to the Paths | 81

Understanding the Default BGP Routing Policy on Packet Transport Routers (PTX Series) | 82

Per-Flow and Per-Prefix Load Balancing Overview | 85

Configuring Per-Prefix Load Balancing | 85

Configuring Per-Flow Load Balancing Based on Hash Values | 87

Configuring Load Balancing Based on MAC Addresses | 88

Load Balancing VPLS Non-Unicast Traffic Across Member Links of an Aggregate Interface | 89

Example: Configuring Multicast Load Balancing over Aggregated Ethernet Links | 90

Configuring Load Balancing for Ethernet Pseudowires

You can configure load balancing for IPv4 traffic over Layer 2 Ethernet pseudowires. You can also configure load balancing for Ethernet pseudowires based on IP information. The option to include IP information in the hash key provides support for Ethernet circuit cross-connect (CCC) connections.

NOTE: This feature is supported only on M120, M320, MX Series, and T Series routers.

To configure load balancing for IPv4 traffic over Layer 2 Ethernet pseudowires, include the `ether-pseudowire` statement at the `[edit forwarding-options hash-key family mpls payload]` hierarchy level:

```
[edit forwarding-options]
hash-key {
    family mpls {
        (label-1 | no-labels);
        payload {
            ether-pseudowire;
        }
    }
}
```

NOTE: You must also configure either the `label-1` or the `no-labels` statement at the `[edit forwarding-options hash-key family mpls]` hierarchy level.

You can also configure load balancing for Ethernet pseudowires based on IP information. This functionality provides support for load balancing for Ethernet cross-circuit connect (CCC) connections. To include IP information in the hash key, include the `ip` statement at the `[edit forwarding-options hash-key family mpls payload]` hierarchy level:

```
[edit forwarding-options]
hash-key {
    family mpls {
        (label-1 | no-labels);
        payload {
```

```
            ip;
        }
    }
}
```

NOTE: You must also configure either the `label-1` or `no-labels` statement at the `[edit forwarding-options hash-key family mpls]` hierarchy level.

You can configure load balancing for IPv4 traffic over Ethernet pseudowires to include only Layer 3 IP information in the hash key. To include only Layer 3 IP information, include the `layer-3-only` option at the `[edit forwarding-options family mpls hash-key payload ip]` hierarchy level:

```
[edit forwarding-options]
hash-key {
    family mpls {
        (label-1 | no-labels);
        payload {
            ip {
                layer-3-only;
            }
        }
    }
}
```

NOTE: You must also configure either the `label-1` or `no-labels` statement at the `[edit forwarding-options hash-key family mpls]` hierarchy level.

Configuring Load-Balance Groups

In addition to including policers in firewall filters, you can configure a load-balance group that is not part of a firewall filter configuration. A load-balance group contains interfaces that all use the same next-hop group characteristic to load-balance the traffic.

To configure a load-balance group, include the `load-balance-group` statement at the `[edit firewall]` hierarchy level:

```
[edit firewall]
load-balance-group group-name {
    next-hop-group[ group-names ];
}
```

Next-hop groups allow you to include multiple interfaces used to forward duplicate packets used in port mirroring. For more information about next-hop groups, see "Configuring Next-Hop Groups to Use Multiple Interfaces to Forward Packets Used in Port Mirroring" on page 42.

Understanding the Algorithm Used to Load Balance Traffic on MX Series Routers

IN THIS SECTION

- MX MPC and T-Series Type 5 FPC Specifics | **63**
- Hashing Algorithm Used in Junos 18.3R1 and later | **63**
- Hash fields used for GRE traffic sent over IPv4 | **64**
- Hash fields used for GRE traffic sent over IPv6 | **66**
- Hash fields used for IPv4 | **67**
- Hash fields used for IPv6 | **69**
- Hash fields used for multiservice | **70**
- Hash fields used for MPLS, Junos 18.3 and later | **72**
- Hash fields used for MPLS from Junos 14.1 to Junos 18.3 | **73**
- List of Junos Updates for Hash Calculation and Load Balancing for MX series routers with MPCs | **74**

When a packet is received on the ingress interface of a device, the packet forwarding engine (PFE) performs a look up to identify the forwarding next hop. If there are multiple equal-cost paths (ECMPs) to the same next-hop destination, the ingress PFE can be configured to distribute the flow between the next hops. Likewise, distribution of traffic may be required between the member links of an aggregated

interface such as aggregated Ethernet. The selection of the actual forwarding next-hop is based on the hash computation result over select packet header fields and several internal fields such as **interface index**. You can configure some of the fields that are used by the hashing algorithm.

- For MX series routers with Modular Port Concentrators (MPCs) and Type 5 FPCs, configure the hash for the supported traffic types at the `forwarding-options enhanced-hash-key` hierarchy level. Details on which fields are included by default for which traffic family can be found below.

  In Junos OS Release 18.3R1, the default method for calculating the enhanced-hash was changed to provide improved entropy for IP tunnels, IPv6 flows and PPPoE payloads transmitted as family multiservice. These defaults can be disabled by setting their respective no- commands.

- For MX series routers with DPCs, configure the hash for the supported traffic types at the `forwarding-options hash-key` hierarchy level.

Junos supports different types of load balancing.

- *Per-prefix load balancing* –Each prefix is mapped to only one forwarding next-hop.

- *Per-packet load balancing*–All next-hop addresses for a destination in the active route are installed in the forwarding table (the term *per-packet* load balancing in Junos is equivalent to what other vendors may call *per-flow* load balancing). See "Configuring Per-Packet Load Balancing" on page 77 for more information.

- *Random packet load balancing*–Next-hops are picked randomly for each packet. This method is available on MX routers with MPC line cards for Aggregated Ethernet interfaces and ECMP paths. To configure per-packet random spray load balancing, include the `per-packet` statement at the `[edit interfaces aex aggregated-ether-options load-balance]` hierarchy level. See *Example: Configuring Aggregated Ethernet Load Balancing* for more information.

- *Per-Packet Random Spray Load Balancing* –When the adaptive load-balancing option fails, per-packet random spray load balancing serves as a last resort. It ensures that the members of ECMP are equally loaded without taking bandwidth into consideration. Per packet causes packet reordering and hence is recommended only if the applications absorb reordering. Per-packet random spray eliminates traffic imbalance that occurs as a result of software errors, except for packet hash.

  Starting in Junos OS Release 20.2R1, you can configure per packet random load balancing on MX240, MX480, and MX960 routers with MPC10E (MPC10E-15C-MRATE and MPC10E-10C-MRATE) line card and MX2010 and MX2020 routers with MX2K-MPC11E line card.

Several additional configuration options are also available:

- *Per-slot hash function configuration* –This method is based on a unique, load-balance hash value for each PIC slot and is only valid for M120, M320, and MX Series routers with DPCE and MS-DPC line cards.

- *Symmetrical load balancing* –This method provides symmetrical load balancing on an 802.3ad LAG. The hash used for symmetrical load balancing is set at the `interface` level of the hierarchy. It ensures that a given flow of duplex traffic traverses the same devices in both directions, and is available on MX Series routers.

MX MPC and T-Series Type 5 FPC Specifics

The hash computation algorithm on MX MPC and T Series Type 5 FPCs produces identical results for packets with swapped layer 3 addresses or layer 4 transport ports. For example, the hash computation result for a packet with source address 192.0.2.1 and destination address 203.0.113.1 is identical to the hash computation result for a packet with source address 203.0.113.1 and destination address 192.0.2.1.

To avoid possible packet re-ordering, layer 4 transport protocol ports are never used in hash computation for fragmented IPv4 packets. This is true for the first fragment of the flow, identified by the `more fragment` bit in a header, and all subsequent fragments, identified by non-zero fragment offset. The first fragment and subsequent fragments are always forwarded over same next-hop.

Hashing Algorithm Used in Junos 18.3R1 and later

In most cases, including layer 3 and layer 4 field information in the hash calculation produces results that are good enough for equitable distribution for traffic. However, in cases such as IP-in-IP or GRE tunneling, layer 3 and layer 4 field information alone may not be enough to produce a hash with sufficient entropy for load balancing. For example, in a deployment where MX series routers transit GRE flows, the GRE encapsulation tunnels typically occur as a single flow with the same source and destination, and same GRE key. Fat flows can also markedly increase the imbalance in link utilization, as traffic volume over the tunnels increases. Another example is when MX PE routers are being used as VPLS PE devices in a subscriber edge deployment where the routers back-haul broadband subscriber traffic from the access devices to a central broadband network gateway (BNG). In such a case, only the subscriber MAC addresses and the BNG router MAC addresses are available for hashing. But with few BNG MACs and relatively few subscriber MACs, the typical layer 3 and layer 4 fields are not sufficient to create a hash for optimal load balancing.

Therefore, for MX series routers with Trio MPCs and running Junos OS Release 18.3R1 or later, the default `enhanced-hash-key` calculation has changed. A summary of the changes is listed here:

- For GRE packets, if the outer IP packet is not a fragmented packet (first fragment or any subsequent fragment), and the inner packet is IPv4 or IPv6, then the source and destination addresses from the inner packet are used in the hash computation in addition to the outer source and destination addresses. Layer 4 ports of the inner packet are also included if the protocol of the inner IP packet is

TCP or UDP, and the inner IP packet is not a fragment (first fragment or any subsequent fragment). Likewise, if the outer IP packet is not a fragment packet, and the inner packet is MPLS, then the top inner label is included in the hash computation.

- For PPPoE packets, if the inner packet is IPv4 or IPv6, then the source and destination addresses from the inner packet are included. Layer 4 ports are included if the protocol of the inner IP packet is TCP or UDP, and the inner IP packet is not a fragment. Inclusion of the PPPoE inner packet fields can be disabled by configuring the `no-payload` option at the `forwarding-options enhanced-hash-key family multiservice` hierarchy level.
- For IPv6, the IPv6 header flow label field is included in the hash computation. RFC 6437 describes the 20-bit flow label field in the IPv6 header. Set the `no-flow-label` option at the `forwarding-options enhanced-hash-key family inet6` hierarchy to disable the new default.

Hash fields used for GRE traffic sent over IPv4

The lists show the fields used in the hash calculation, for non-fragmented packets, in Junos 18.3R1 and later. By default, the field is used in the hash calculation unless otherwise noted. Also where noted, the IP and port fields used in the hash is symmetric, that is, swapping the fields does not change the hash result.

- **IPv4, GRE**
  - GRE Key
  - Source and destination address; symmetric
  - Protocol
  - DSCP (disabled)
  - Incoming Interface Index (disabled)
- **IPv4 in IPv4, GRE**
  - Payload (inner IPv4: source and destination ports, IP addresses); symmetric
  - GRE Key

    GRE Protocol = IPv4
  - Source and destination address; symmetric
  - Protocol

- DSCP (disabled)
- Incoming Interface Index (disabled)

- **IPv6 in IPv4, GRE**
  - Payload (inner IPv6: source and destination ports, IP addresses); symmetric
  - GRE Key

    GRE Protocol = IPv6
  - Source and destination address; symmetric
  - Protocol
  - DSCP (disabled)
  - Incoming Interface Index (disabled)
- **MPLS in IPv4, GRE**
  - Payload (inner MPLS: top label)
  - GRE Key

    GRE Protocol = MPLS
  - Source and destination address; symmetric
  - Protocol
  - DSCP (disabled)
  - Incoming Interface Index (disabled)
- **IPv4, L2TPv2 used in Junos 17.2 and later**

  Inclusion of the L2TPv2 tunnel ID and session ID can be enabled by configuring the `forwarding-options enhanced-hash-key family inet l2tp-tunnel-session-identifier` option. Note that Juniper does not recommend enabling this option by default. This is because L2TP session identification is based on the destination UDP port match (1701), and this port may not be exclusively used for L2TP transport so the extraction of the tunnel and session ID fields from the packet may not always be accurate.
  - Session ID
  - Tunnel ID
  - Source and destination port
  - Source and destination address; symmetric

- Protocol (UDP)
- DSCP (disabled)
- Incoming Interface Index (disabled)

Hash fields used for GRE traffic sent over IPv6

The list shows the fields used in the hash calculation for non-fragmented packets. By default, the field is used in the hash calculation unless otherwise noted. Also where noted, the IP and port fields used in the hash is symmetric, that is, swapping the fields does not change the hash result.

- **IPv6, GRE**
  - GRE Key
  - Source and destination address; symmetric
  - Next header
  - Flow label (Junos 18.3 and later)
  - Traffic class (disabled)
  - Incoming Interface Index (disabled)
- **IPv4 in IPv6, GRE (Junos 18.3 and later)**
  - Payload (inner IPv4: source and destination ports, IP addresses); symmetric
  - GRE Key

    GRE Protocol = IPv4
  - Source and destination address; symmetric
  - Next header
  - Flow label (Junos 18.3 and later)
  - Traffic class (disabled)
  - Incoming Interface Index (disabled)
- **IPv6 in IPv6, GRE (Junos 18.3 and later)**
  - Payload (inner IPv6: source and destination ports, IP addresses); symmetric

- GRE Key

  GRE Protocol = IPv6

- Source and destination address; symmetric
- Next header
- Flow label (Junos 18.3 and later)
- Traffic class (disabled)
- Incoming Interface Index (disabled)

- **MPLS in IPv6, GRE (Junos 18.3 and later)**
  - Payload (inner MPLS: top labels); symmetric
  - GRE Key

    GRE Protocol = MPLS

  - Source and destination address; symmetric
  - Next header
  - Flow label
  - Traffic class (disabled)
  - Incoming Interface Index (disabled)

Hash fields used for IPv4

The list shows the fields used in the hash calculation for non-fragmented packets, except where noted. By default, the field is used in the hash calculation unless otherwise noted. Also where noted, the IP and port fields hash is symmetric, that is, swapping the fields does not change the hash result.

- **IPv4, not TCP or UDP, or fragmented packets**
  - Source and destination address; symmetric
  - Protocol
  - DSCP (disabled)
  - Incoming Interface Index (disabled)

- **IPv4, TCP and UDP, non fragmented packets**
  - Source and destination port; symmetric
  - Source and destination address; symmetric
  - Protocol
  - DSCP (disabled)
  - Incoming Interface Index (disabled)
- **IPv4, PPTP**
  - 16 least significant bits of the GRE Key
  - Source and destination address; symmetric
  - Protocol
  - DSCP (disabled)
  - Incoming Interface Index (disabled)
- **IPv4, GTP, UDP traffic to destination port 2152**

  Inclusion of GPRS tunneling protocol (GTP) tunnel endpoint identifier (TEID) can be enabled at the `forwarding-options enhanced-hash-key family inet gtp-tunnel-endpoint-identifier` option. Note that Juniper does not recommend enabling this option by default. This is because GTP session identification is based on the destination UDP port match (2152), and this port may not be exclusively used for GTP transport, so the extraction of TEID field from the packet may not always be accurate.

  - GTP TEID (disabled)
  - Source and destination port
  - Source and destination address; symmetric
  - Protocol
  - DSCP (disabled)
  - Incoming Interface Index (disabled)

Hash fields used for IPv6

The list shows the fields used in the hash calculation for non-fragmented packets, except where noted. By default, the field is used in the hash calculation unless otherwise noted. Also where noted, the IP and port fields hash is symmetric, that is, swapping the fields does not change the hash result.

- **IPv6, non TCP and UDP packet, or TCP and UDP packet fragmented by the originator**
  - Source and destination address; symmetric
  - Next header
  - Flow label (Junos 18.3 and later)
  - Traffic class (disabled)
  - Incoming Interface Index (disabled)
- **IPv6, non fragmented TCP and UDP packet**
  - Source and destination port; symmetric
  - Source and destination address; symmetric
  - Next header
  - Flow label (Junos 18.3 and later)
  - Traffic class (disabled)
  - Incoming Interface Index (disabled)
- **IPv6, PPTP**
  - 16 least significant bits of the GRE Key
  - Source and destination address; symmetric
  - Next header
  - Flow label (Junos 18.3 and later)
  - Traffic class (disabled)
  - Incoming Interface Index (disabled)
- **IPv6, GTP**

  Inclusion of GPRS tunneling protocol (GTP) tunnel endpoint identifier (TEID) can be enabled at the `forwarding-options enhanced-hash-key family inet gtp-tunnel-endpoint-identifier` hierarchy level. Note that

Juniper does not recommend enabling this option by default. This is because GTP session identification is based on the destination UDP port match (2152), and this port may not be exclusively used for GTP transport, so the extraction of TEID field from the packet may not always be accurate.

- GTP TEID (disabled by default; enable at the `forwarding-options enhanced-hash-key family inet gtp-tunnel-endpoint-identifier` hierarchy level.
- Source and destination port
- Source and destination address; symmetric
- Next header
- Flow label (Junos 18.3 and later)
- Traffic class (disabled)
- Incoming Interface Index (disabled)

Hash fields used for multiservice

Family multiservice hash configuration applies to packets entering into the router as `family ccc`, `vpls`, or `bridge`. The list shows the fields used in the hash calculation for non-fragmented packets. By default, the field is used in the hash calculation unless otherwise noted. Also where noted, the IP and port fields used in the hash is symmetric, that is, swapping the fields does not change the hash result.

- **Ethernet, non-IP or non-MPLS**

  If configured, payload information is extracted from untagged packets or packets with up to two VLAN tags.

  - Outer 802.1p (disabled)
  - Source and destination MAC; symmetric
  - Incoming Interface Index (disabled)

- **Ethernet, IPv4**

  - Payload (inner IPv4: source and destination ports, IP addresses); symmetric
  - Outer 802.1p (disabled)
  - Source and destination MAC; symmetric

    - Incoming Interface Index (disabled)

- **Ethernet, IPv6**
    - Payload (inner IPv6: source and destination ports, IP addresses); symmetric
    - Outer 802.1p (disabled)
    - Source and destination MAC; symmetric
    - Incoming Interface Index (disabled)
- **Ethernet, MPLS**
    - Payload (inner MPLS: top labels plus inner IPv4 and IPv6 fields); symmetric. See *Hash fields used for MPLS, Junos 18.3 and later*, below, for related information.
    - Outer 802.1p (disabled)
    - Source and destination MAC; symmetric
    - Incoming Interface Index (disabled)
- **IPv4 in PPPoE (data packet)**
    - Payload (inner IPv4: source and destination ports, IP addresses); symmetric
    - PPP protocol IPv4 version 0x1, type 0x1
    - Outer 802.1p (disabled)
    - Source and destination MAC; symmetric
    - Incoming Interface Index (disabled)
- **IPv6 in PPPoE (data packet)**
    - Payload (inner IPv6: source and destination ports, IP addresses); symmetric
    - PPP protocol IPv6 version 0x1, type 0x1
    - Outer 802.1p (disabled)
    - Source and destination MAC; symmetric
    - Incoming Interface Index (disabled)

Hash fields used for MPLS, Junos 18.3 and later

The list shows the fields used in the hash calculation for non-fragmented packets. By default, the field is used in the hash calculation unless otherwise noted. Also where noted, the IP and port fields used in the hash is symmetric, that is, swapping the fields does not change the hash result.

- **MPLS, Encapsulated IPv4 or IPv6**
  - Payload (inner IPv4: source and destination ports, IP addresses); symmetric
  - Payload (inner IPv6: source and destination ports, IP addresses, next header); symmetric
  - Label 1..16 (20 bits)
  - Outer Label EXP (disabled)
  - Incoming Interface Index (disabled)
- **MPLS, IPv4 or IPv6 in Ethernet pseudo-wire**
  - Payload (IPv4/IPv6 in Ethernet pseudo-wire)
  - Label 2..16 (20 bits)
  - Outer Label EXP (disabled)
  - Label 1 (20 bits)
  - Incoming Interface Index (disabled)
- **MPLS, MPLS in Ethernet pseudo-wire**
  - Payload (two top labels of MPLS label stack entry in Ethernet pseudo-wire)
  - Label 2..16 (20 bits)
  - Outer Label EXP (disabled)
  - Label 1 (20 bits)
  - Incoming Interface Index (disabled)
- **MPLS, entropy label**

  When an entropy label is detected, the payload field is not processed, and the indicator is not included into hash computation

  - Label 1..16 (20 bits)
  - Outer Label EXP (disabled)

- Incoming Interface Index (disabled)

Hash fields used for MPLS from Junos 14.1 to Junos 18.3

The list shows the fields used in the hash calculation for non-fragmented packets. By default, the field is used in the hash calculation unless otherwise noted. Also where noted, the IP and port fields used in the hash is symmetric, that is, swapping the fields does not change the hash result.

- **MPLS, Encapsulated IPv4 or IPv6**
  - Payload (inner IPv4: source and destination ports, IP addresses); symmetric

    Payload (inner IPv6: source and destination ports, IP addresses, next header); symmetric
  - Label 2.8 (20 bits)

    Outer Label EXP (disabled)

    Label 1 (20 bits)
  - Incoming Interface Index (disabled)
- **MPLS, IPv4 or IPv6 in Ethernet pseudo-wire**
  - Payload (IPv4/IPv6 in Ethernet pseudo-wire)
  - Label 2.8 (20 bits)

    Outer Label EXP (disabled)

    Label 1 (20 bits)
  - Incoming Interface Index (disabled)
- **MPLS, MPLS in Ethernet pseudo-wire**
  - Payload (two top labels of MPLS label stack entry in Ethernet pseudo-wire)
  - Label 2..16 (20 bits)
  - Outer Label EXP (disabled)
  - Label 1 (20 bits)
  - Incoming Interface Index (disabled)
- **MPLS, entropy label**

When an entropy label is detected, the payload field is not processed, and the indicator is not included into hash computation

- Label 2.8 (20 bits)

  Outer Label EXP (disabled)

  Label 1 (20 bits)

- Incoming Interface Index (disabled)

List of Junos Updates for Hash Calculation and Load Balancing for MX series routers with MPCs

Table 1: List of updates for MX series routers

Junos Release	Change
18.3R1	Includes IPv6 flow label, inner GRE header, and inner PPPoE in default hash computation. Increases MPLS label stack depth to 16 labels.
17.2R1	Load balancing for L2TP encapsulated IPv4 and IPv6 packets.
16.1R1	Includes EoMPLS payload hash with control word. Introduces source-only and destination-only based hashing.
15.1R1	Provides targeted distribution of static interfaces across AE member links. Includes source, destination, and MAC of MPLS encapsulated PPPoE payload in the default hash computation.
14.2R3	Increases scaling of LAG and MC-LAG.
14.2R2	Provides aggregate Ethernet bundle with 10G, 40G and 100G links.

14.1R1	Decouples aeX interface creation from `ch agg eth dev`. Increases aggregate Ethernet interface name space. Provides adaptive load balancing for ECMP next hops.
13.3R1	Includes enhancements for adaptive, per-packet-random, and periodic-rebalance load balancing.
11.4R1	provides load sharing across ECMP next hops.

Understanding Per-Packet Load Balancing

By default, when there are multiple equal-cost paths to the same destination for the active route, Junos OS uses a hash algorithm to choose one of the next-hop addresses to install in the forwarding table. Whenever the set of next hops for a destination changes in any way, the next-hop address is re-chosen using the hash algorithm. Starting in Junos OS Release 18.3R1, for MX series routers, the default behavior for IPv6, GRE, and PPPoE packet hash computation was modified to include the flow-label field for improved load-balancing in certain cases (you can use the `no-payload` option to revert to the previous method for hash computation). See "Understanding the Algorithm Used to Load Balance Traffic on MX Series Routers" on page 61 for details.

You can configure Junos OS so that, for the active route, all next-hop addresses for a destination are installed in the forwarding table. This feature is called *per-packet load balancing*. The naming may be counter-intuitive. However, Junos *per-packet* load balancing is functionally equivalent to what other vendors may term *per-flow* load balancing. You can use load balancing to spread traffic across multiple paths between routers.

Figure 2 on page 76 shows a simple load balancing scenario. Device R1 is in AS 64500 and is connected to both Device R2 and Device R3, which are in AS 64501. Device R1 can be configured to load balance traffic across the two links.

Figure 2: Simple Load Balancing Scenario

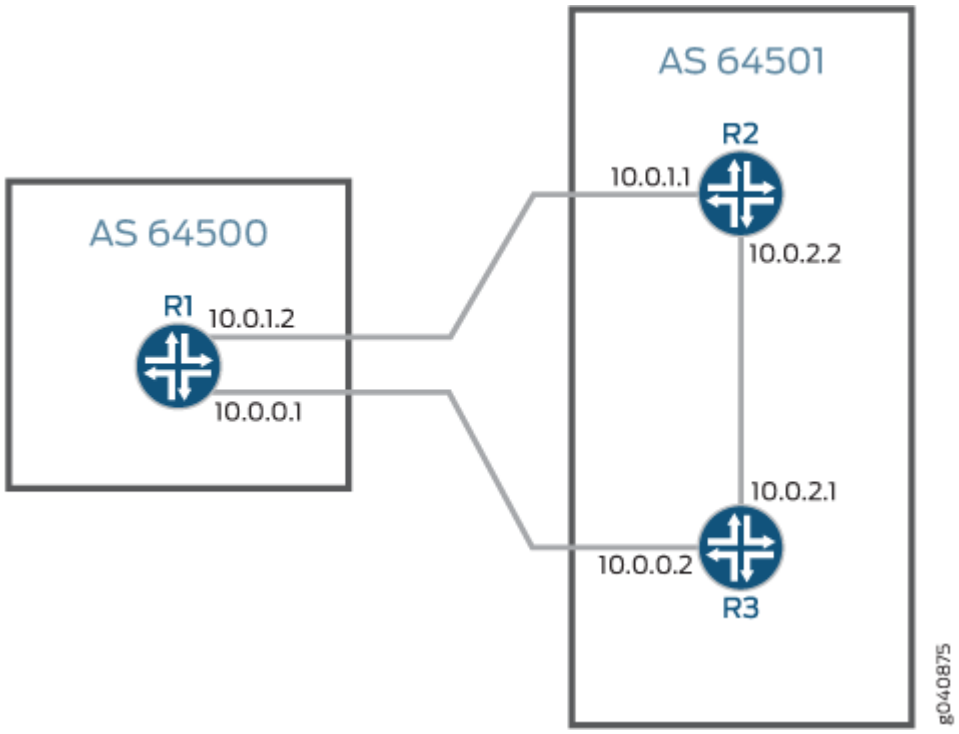

Starting in Junos OS 13.3R3, for MX Series 5G Universal Routing Platforms with modular port concentrators (MPCs) only, you can configure consistent load balancing, which prevents the reordering of all flows to active paths in an equal-cost multipath (ECMP) group when one or more next-hop paths fail. Only flows for paths that are inactive are redirected to another active next-hop path. Flows mapped to servers that remain active are maintained. This feature applies only to external BGP peers.

Starting in Junos OS Release 19.1R1,on QFX10000 switches, you can configure load balancing of IPv4 or IPv6 packets by using GPRS Tunneling Protocol-tunnel endpoint identifier (GTP-TEID) field hash calculations. The GTP-TEID hashing is added to the Layer 2 and Layer 3 field hashing that you have already configured. To enable this feature on QFX10000 switches, configure the `gtp-tunnel-endpoint-identifier` statement at the `[edit forwarding-options enhanced-hash-key family inet]` or the `[edit forwarding-options enhanced-hash-key family inet6]` hierarchy Level. GTP versions 1 and 2 are supported; they support only user data. You must use UDP port number 2152 for both GTP versions.

Release History Table

Release	Description
19.1R1	Starting in Junos OS Release 19.1R1,on QFX10000 switches, you can configure load balancing of IPv4 or IPv6 packets by using GPRS Tunneling Protocol-tunnel endpoint identifier (GTP-TEID) field hash calculations

18.3R1	Starting in Junos OS Release 18.3R1, for MX series routers, the default behavior for IPv6, GRE, and PPPoE packet hash computation was modified to include the flow-label field for improved load-balancing in certain cases (you can use the `no-payload` option to revert to the previous method for hash computation).
13.3R3	Starting in Junos OS 13.3R3, for MX Series 5G Universal Routing Platforms with modular port concentrators (MPCs) only, you can configure consistent load balancing, which prevents the reordering of all flows to active paths in an equal-cost multipath (ECMP) group when one or more next-hop paths fail.

RELATED DOCUMENTATION

- Example: Load Balancing BGP Traffic
- Configuring Per-Packet Load Balancing | **77**
- Configuring Load Balancing Based on MPLS Labels
- *Configuring Load Balancing for Ethernet Pseudowires*
- *Configuring Load Balancing Based on MAC Addresses*
- *Configuring VPLS Load Balancing Based on IP and MPLS Information*
- *Configuring VPLS Load Balancing on MX Series 5G Universal Routing Platforms*
- Configuring Consistent Load Balancing for ECMP Groups

Configuring Per-Packet Load Balancing

IN THIS SECTION

- Per-Packet Load Balancing Examples | **80**

In Junos OS, you enable per-flow load balancing by setting the **load-balance per-packet** action in the routing policy configuration. The naming may be counter-intuitive, because in Junos, *per-packet* load balancing is functionally equivalent to what other vendors may term *per-flow* load balancing.

To configure per-packet load balancing, include the `load-balance per-packet` statement either as an option of the `route-filter` statement at the `[edit policy-options policy-statement` *`policy-name`* `term` *`term-name`* `from]` hierarchy level:

```
[edit policy-options policy-statement policy-name term term-name from]
route-filter destination-prefix match-type {
    load-balance per-packet;
}
```

or at the `[edit policy-options policy-statement` *`policy-name`* `term` *`term-name`* `then]` hierarchy level:

```
[edit policy-options policy-statement policy-name term term-name then]
load-balance per-packet;
```

To complete the configuration you must apply the routing policy to routes exported from the routing table to the forwarding table, by including the policy name in the list specified by the `export` statement:

```
export [ policy-names ];
```

You can include this statement at the following hierarchy levels:

- `[edit routing-options forwarding-table]`
- `[edit logical-systems` *`logical-system-name`* `routing-options forwarding-table]`

By default, Junos ignores port data when determining flows. To include port data in the flow determination, include the `family inet` statement at the `[edit forwarding-options hash-key]` hierarchy level:

```
[edit forwarding-options hash-key]
family inet {
    layer-3;
    layer-4;
}
```

If you include both the **layer 3** and **layer 4** statements, the device uses the following Layer 3 and Layer 4 information to load-balance:

- Source IP address
- Destination IP address
- Protocol

- Source port number
- Destination port number
- Incoming interface index
- IP type of service

When all of the **layer 3** and **layer 4** parameters are identical, the device sends packets in the flow through the same interface, which in turn helps prevent out of order delivery for TCP and UDP flows..

Internet Control Message Protocol (ICMP) packets are handled differently because the field location offset is the checksum field, which makes each ping packet a separate “flow.” There are other protocols that can be encapsulated in IP that may have a varying value in the 32-bit offset. This may also be problematic because these protocols are seen as a separate flow.

With M Series (with the exception of the M120 router) and T Series routers, the first fragment is mapped to the same load-balanced destination as the unfragmented packets. The other fragments can be mapped to other load-balanced destinations.

For the M120 router only, all fragments are mapped to the same load-balanced destination. This destination is not necessarily the same as that for unfragmented packets.

By default, or if you include only the **layer 3** statement, the router uses the incoming interface index as well as the following Layer 3 information in the packet header to load balance traffic:

- Source IP address
- Destination IP address
- Protocol

By default, IP version 6 (IPv6) packets are automatically load-balanced based on the following Layer 3 and Layer 4 information:

- Source IP address
- Destination IP address
- Protocol
- Source port number
- Destination port number
- Incoming interface index
- Traffic class

Per-Packet Load Balancing Examples

Perform per-packet load balancing for all routes:

```
[edit]
policy-options {
    policy-statement load-balancing-policy {
        then {
            load-balance per-packet;
        }
    }
}
routing-options {
    forwarding-table {
        export load-balancing-policy;
    }
}
```

Perform per-packet load balancing only for a limited set of routes:

```
[edit]
policy-options {
    policy-statement load-balancing-policy {
        from {
            route-filter 192.168.10/24 orlonger;
            route-filter 10.114/16 orlonger;
        }
        then {
            load-balance per-packet;
        }
    }
}
routing-options {
    forwarding-table {
        export load-balancing-policy;
    }
}
```

To configure per-packet random spray load balancing, include the `load-balance random` statement at the `[edit policy-options policy-statement` *`policy-name`* `term` *`term-name`* `then]` hierarchy level:

```
[edit policy-options policy-statement policy-name term term-name then]
load-balance random;
```

To complete the configuration you must apply the routing policy to routes exported from the routing table to the forwarding table, by including the policy name in the list specified by the `export` statement at the `[edit routing-options forwarding-table]` hierarchy level

```
[edit routing-options forwarding-table]
export [ policy-names ];
```

RELATED DOCUMENTATION

Understanding Per-Packet Load Balancing | 75

Understanding Load Balancing for BGP Traffic with Unequal Bandwidth Allocated to the Paths

The multipath option removes the tiebreakers from the active route decision process, thereby allowing otherwise equal cost BGP routes learned from multiple sources to be installed into the forwarding table. However, when the available paths are not equal cost, you may wish to load balance the traffic asymmetrically.

Once multiple next hops are installed in the forwarding table, a specific forwarding next hop is selected by the Junos OS per-prefix load-balancing algorithm. This process hashes against a packet's source and destination addresses to deterministically map the prefix pairing onto one of the available next hops. Per-prefix mapping works best when the hash function is presented with a large number of prefixes, such as might occur on an Internet peering exchange, and it serves to prevent packet reordering among pairs of communicating nodes.

An enterprise network normally wants to alter the default behavior to evoke a *per-packet* load-balancing algorithm. Per-packet is emphasized here because its use is a misnomer that stems from the historic behavior of the original Internet Processor ASIC. In reality, current Juniper Networks routers support per-prefix (default) and per-flow load balancing. The latter involves hashing against various Layer 3 and

Layer 4 headers, including portions of the source address, destination address, transport protocol, incoming interface, and application ports. The effect is that now individual flows are hashed to a specific next hop, resulting in a more even distribution across available next hops, especially when routing between fewer source and destination pairs.

With per-packet load balancing, packets comprising a communication stream between two endpoints might be resequenced, but packets within individual flows maintain correct sequencing. Whether you opt for per-prefix or per-packet load balancing, asymmetry of access links can present a technical challenge. Either way, the prefixes or flows that are mapped to, for example, a T1 link will exhibit degraded performance when compared to those flows that map to, for example, a Fast Ethernet access link. Worse yet, with heavy traffic loads, any attempt at equal load balancing is likely to result in total saturation of the T1 link and session disruption stemming from packet loss.

Fortunately, the Juniper Networks BGP implementation supports the notion of a bandwidth community. This extended community encodes the bandwidth of a given next hop, and when combined with multipath, the load-balancing algorithm distributes flows across the set of next hops proportional to their relative bandwidths. Put another way, if you have a 10-Mbps and a 1-Mbps next hop, on average nine flows will map to the high-speed next hop for every one that uses the low speed.

Use of BGP bandwidth community is supported only with per-packet load balancing.

The configuration task has two parts:

- Configure the external BGP (EBGP) peering sessions, enable multipath, and define an import policy to tag routes with a bandwidth community that reflects link speed.
- Enable per-packet (really per-flow) load balancing for optimal distribution of traffic.

RELATED DOCUMENTATION

Understanding Per-Packet Load Balancing | 75

Understanding the Default BGP Routing Policy on Packet Transport Routers (PTX Series)

On PTX Series Packet Transport Routers, the default BGP routing policy differs from that of other Junos OS routing devices.

The PTX Series routers are MPLS transit platforms that do IP forwarding, typically using interior gateway protocol (IGP) routes. The PTX Series Packet Forwarding Engine can accommodate a relatively small number of variable-length prefixes.

NOTE: A PTX Series router can support full BGP routes in the control plane so that it can be used as a route reflector (RR). It can do exact-length lookup multicast forwarding and can build the multicast forwarding plane for use by the unicast control plane (for example. to perform a reverse-path forwarding lookup for multicast).

Given the PFE limitation, the default routing policy for PTX Series routers is for BGP routes not to be installed in the forwarding table. You can override the default routing policy and select certain BGP routes to install in the forwarding table.

The default behavior for load balancing and BGP routes on PTX Series routers is as follows. It has the following desirable characteristics:

- Allows you to override the default behavior without needing to alter the default policy directly
- Reduces the chance of accidental changes that nullify the defaults
- Sets no flow-control actions, such as accept and reject

The default routing policy on the PTX Series routers is as follows:

```
user@host# show policy-options | display inheritance defaults no-comments
policy-options {
    policy-statement junos-ptx-series-default {
        term t1 {
            from {
                protocol bgp;
                rib inet.0;
            }
            then no-install-to-fib;
        }
        term t2 {
            from {
                protocol bgp;
                rib inet6.0;
            }
            then no-install-to-fib;
        }
        term t3 {
```

```
            then load-balance per-packet;
        }
    }
}
routing-options {
    forwarding-table {
        default-export junos-ptx-series-default;
    }
}
user@host# show routing-options forwarding-table default-export | display inheritance defaults
no-comments
default-export junos-ptx-series-default;
```

As shown here, the `junos-ptx-series-default` policy is defined in `[edit policy-options]`. The policy is applied in `[edit routing-options forwarding-table]`, using the `default-export` statement. You can view these default configurations by using the `| display inheritance` flag.

Also, you can use the `show policy` command to view the default policy.

```
user@host> show policy junos-ptx-series-default
Policy junos-ptx-series-default:
    Term t1:
        from proto BGP
         inet.0
        then install-to-fib no
    Term t2:
        from proto BGP
         inet6.0
        then install-to-fib no
    Term t3:
        then load-balance per-packet
```

CAUTION: We strongly recommend that you do not alter the `junos-ptx-series-default` routing policy directly.

Junos OS chains the `junos-ptx-series-default` policy and any user-configured export policy. Because the `junos-ptx-series-default` policy does not use flow-control actions, any export policy that you configure is executed (by way of the implicit next-policy action) for every route. Thus you can override any actions set by the `junos-ptx-series-default` policy. If you do not configure an export policy, the actions set by `junos-ptx-series-default` policy are the only actions.

You can use the policy action `install-to-fib` to override the `no-install-to-fib` action.

Similarly, you can set the `load-balance per-prefix` action to override the `load-balance per-packet` action.

RELATED DOCUMENTATION

> *Conditional Advertisement and Import Policy (Routing Table) with certain match conditions*

Per-Flow and Per-Prefix Load Balancing Overview

By default, when there are multiple equal-cost paths to the same destination, Junos OS chooses one of the next-hop addresses at random.

On all M Series Multiservice Edge Routers, MX Series 5G Universal Routing Platforms, and T Series Core Routers, you have the additional option of configuring per-prefix load balancing based on a specified hash value that enables the router to elect a next hop independently of the route chosen by other routers.

On the M120, M320, and MX Series routers only, you have the additional option of enabling per-flow load balancing based on a unique, load-balance hash value for each Packet Forwarding Engine slot.

RELATED DOCUMENTATION

> Configuring Per-Prefix Load Balancing | **85**
>
> Configuring Per-Flow Load Balancing Based on Hash Values | **87**

Configuring Per-Prefix Load Balancing

By default, Junos OS uses a hashing method based only on the destination address to elect a forwarding next hop when multiple equal-cost paths are available. As a result, when multiple routers or switches share the same set of forwarding next hops for a given destination, they can elect the same forwarding next hop.

You can enable router-specific or switch-specific load balancing by including a per-prefix hash value. However, this method applies only to indirect next hops. In other words, when we have a route with a protocol next hop that is not directly connected, it can be resolved over a set of equal-cost forwarding

next hops. Only in this case, we use the hashing algorithm to elect a forwarding next hop. An example of this is routes learned from an IBGP neighbor. The protocol next hop for those routes might not be directly reachable and would be resolved through some IGP or static routes. The result could be a set of equal-cost forwarding next hops to reach that protocol next hop. Per-prefix load balancing thus leads to better utilization of the available links.

To configure per-prefix load balancing, include the `load-balance` statement at the `[edit forwarding-options]` hierarchy level:

```
[edit forwarding-options]
load-balance {
    indexed-load-balance;
    per-prefix {
        hash-seed number;
    }
}
```

To enable per-prefix load balancing, you must include the `hash-seed` *`number`* statement. The range that you can configure is 0 (the default) through 65,535. If no hash seed is configured, the elected forwarding next hop is the same as in previous releases.

If you notice an issue with the load-balance distribution, try including the `indexed-load-balance` statement at the `[edit forwarding-options load-balance]` hierarchy level. This statement causes the creation of a nexthop structure that is both a function of the hash, and a function of the low-order bits of the IP address.

For MPC line cards in MX routers, `indexed-load-balance` has been superseded by an internal hash-rotation mechanism to reduce polarization.

CAUTION: Including the `indexed-load-balance` statement causes an increase in memory usage on the device.

```
indexed-load-balance;
```

RELATED DOCUMENTATION

Understanding the Algorithm Used to Load Balance Traffic on MX Series Routers | 61

enhanced-hash-key | 158

Configuring Per-Flow Load Balancing Based on Hash Values

By default, Junos OS uses a hashing method based only on the destination address to elect a forwarding next hop when multiple equal-cost paths are available. All Packet Forwarding Engine slots are assigned the same hash value by default.

You can enable router-specific or switch-specific load balancing by configuring the router or switch to assign a unique, load-balance hash value for each Packet Forwarding Engine slot.

NOTE: This feature is supported only on M120, M320, and MX Series routers.

To configure per-flow load balancing. include the `load-balance` statement at the `[edit forwarding-options]` hierarchy level:

```
[edit forwarding-options]
load-balance {
    indexed-load-balance;
    per-flow {
        hash-seed;
    }
}
```

To enable per-flow load balancing, you must include the `hash-seed` statement. Junos OS automatically chooses a value for the hashing algorithm. You cannot configure a specific value for the `hash-seed` statement when you enable per-flow load balancing.

Configuring Load Balancing Based on MAC Addresses

The hash key mechanism for load-balancing uses Layer 2 media access control (MAC) information such as frame source and destination address. To load-balance traffic based on Layer 2 MAC information, include the `family multiservice` statement at the `[edit forwarding-options hash-key]` hierarchy level:

```
family multiservice {
    destination-mac;
    source-mac;
}
```

To include the destination-address MAC information in the hash key, include the **destination-mac** option. To include the source-address MAC information in the hash key, include the **source-mac** option.

NOTE: Any packets that have the same source and destination address will be sent over the same path.

NOTE: You can configure per-packet load balancing to optimize VPLS traffic flows across multiple paths.

NOTE: Aggregated Ethernet member links will now use the physical MAC address as the source MAC address in 802.3ah OAM packets.

NOTE: ACX Series routers do not support VPLS.

Load Balancing VPLS Non-Unicast Traffic Across Member Links of an Aggregate Interface

By default, VPLS non-unicast (or BUM — broadcast, unknown, and multicast) traffic sent across aggregate Ethernet interfaces is sent across only one member link of the aggregate interface. Starting with Junos OS Release 14.1, you can configure each VPLS instance to load balance BUM traffic across all members of an aggregate interface. This is referred to as BUM hashing.

To enable BUM hashing for an VPLS instance, add `bum-hashing` to the routing instance at the `[edit routing-instances` *`instance-name`* `protocols vpls]` hierarchy level. For example:

```
[edit routing-instances]
instance-name {
    protocols {
        vpls {
            bum-hashing;
        }
    }
}
```

WARNING: Enabling or disabling BUM hashing on a VPLS routing instance causes the routing instance to be destroyed and re-created when the configuration change is committed.

You can also specify which forwarding class to use for forwarding BUM traffic. When CoS-based forwarding (CBF) is configured on a VPLS PE router, BUM traffic uses the default forwarding class to select the label-switched path (LSP). Starting with Junos OS Release 14.1, you can associate an LSP with the default forwarding class.

To associate an LSP with the default forwarding class, add the `forwarding-class-default` statement at the `[edit class-of-service forwarding-policy next-hop-map` *`next-hop-map-name`*`]` hierarchy level. For example:

```
[edit class-of-service forwarding-policy next-hop-map next-hop-map-name]
forwarding-class-default {
    lsp-next-hop value;
}
```

Release History Table

Release	Description
14.1	Starting with Junos OS Release 14.1, you can configure each VPLS instance to load balance BUM traffic across all members of an aggregate interface.
14.1	Starting with Junos OS Release 14.1, you can associate an LSP with the default forwarding class.

RELATED DOCUMENTATION

bum-hashing | 141

Example: Configuring Multicast Load Balancing over Aggregated Ethernet Links | 90

Example: Configuring Multicast Load Balancing over Aggregated Ethernet Links

IN THIS SECTION

- Requirements | 91
- Overview | 91
- Configuration | 92
- Verification | 107

This example shows how to configure point-to-multipoint LSPs to load balance across aggregated Ethernet links. The load balancing applies to all traffic types, including multicast. Feature parity for multicast load balancing of point-to-multipoint LSPs over aggregated Ethernet child links on the MX Series routers with MPCs or MICs is supported in Junos OS Releases 11.1R2, 11.2R2, and 11.4.

NOTE: VPLS multicast load balancing requires Junos OS Release 14.1 or later.

Requirements

Before you begin:

1. Configure the router interfaces.
2. Configure an interior gateway protocol or static routing. See the Junos OS Routing Protocols Library for Routing Devices.

Overview

IN THIS SECTION

- Topology | 92

This example shows a sample topology and configuration to perform the following tasks:

- Load balancing VPLS multicast traffic over link aggregation
- Load balancing point-to-multipoint multicast traffic over link aggregation
- Re-load balancing after a change in the next-hop topology

  Next-hop topology changes might include but are not limited to:

  - Layer 2 membership change in the link aggregation
  - Indirect next-hop change
  - Composite next-hop change

Load balancing is hash-based, so the higher the number of flows, the better. As is the case with unicast, you can also configure the hash key to be based on Layer 3 and Layer 4 information to achieve a better load-balancing result. There are a few exceptions that are specific to multicast traffic, which might lead to uneven load balancing—for example, when the outgoing interface list includes multiple aggregated interfaces with an unequal number of child links.

NOTE: For Draft Rosen multicast VPNs (MVPNs), load balancing over aggregated Ethernet interfaces is uneven when the LAGs are all core interfaces. In the case of Next-Generation MBGP MPVNs, multicast traffic is sent over point-to-multipoint and RSVP, and the hash is computed up to the IP headers. In the Draft Rosen case, multicast traffic is tunneled over GRE tunnels, and the hash is used only on GRE tunnel headers. This is why load balancing is not even for Draft Rosen when the LAGs are all core interfaces.

Topology

Figure 3 on page 92 shows the topology for this example. The example includes the configuration for Devices PE1 and PE2.

Figure 3: Multicast Load Balancing over Aggregated Ethernet Links

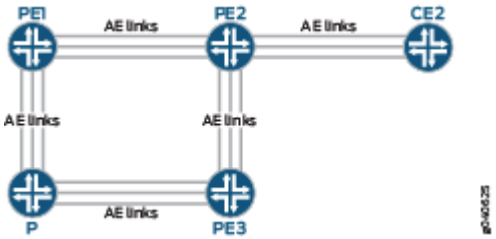

Configuration

IN THIS SECTION

- Procedure | 93
- Results | 99

Procedure

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the `[edit]` hierarchy level.

Device PE1

```
[edit]
set forwarding-options hash-key family multiservice source-mac
set forwarding-options hash-key family multiservice destination-mac
set forwarding-options hash-key family multiservice payload ip layer-3
set interfaces ge-0/0/6 gigether-options 802.3ad ae0
set interfaces ge-0/1/6 gigether-options 802.3ad ae0
set interfaces ge-0/2/2 encapsulation ethernet-vpls
set interfaces ge-0/2/2 unit 0 family vpls
set interfaces ge-0/2/3 gigether-options 802.3ad ae0
set interfaces ge-0/2/6 gigether-options 802.3ad ae0
set interfaces ge-0/3/0 gigether-options 802.3ad ae0
set interfaces ge-0/3/1 gigether-options 802.3ad ae0
set interfaces ge-0/3/6 gigether-options 802.3ad ae0
set interfaces ge-1/0/6 gigether-options 802.3ad ae0
set interfaces ge-1/2/6 unit 0 family inet address 10.13.1.2/30
set interfaces ae0 unit 0 family inet address 10.11.11.1/30
set interfaces ae0 unit 0 family iso
set interfaces ae0 unit 0 family mpls
set policy-options policy-statement exp-to-fwd term a from community grn-com
set policy-options policy-statement exp-to-fwd term a then install-nexthop lsp PE1-to-PE2
set policy-options policy-statement exp-to-fwd term a then accept
set policy-options community grn-com members target:65000:1
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls label-switched-path PE1-to-PE2 to 10.255.19.77
set protocols mpls label-switched-path PE1-to-PE3 to 10.255.19.79
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group int type internal
set protocols bgp group int local-address 10.255.71.214
set protocols bgp group int family inet any
set protocols bgp group int family l2vpn signaling
```

```
set protocols bgp group int neighbor 10.255.19.77
set protocols bgp group int neighbor 10.255.19.79
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set routing-instances vpls instance-type vpls
set routing-instances vpls interface ge-0/2/2.0
set routing-instances vpls route-distinguisher 65000:1
set routing-instances vpls vrf-target target:65000:1
set routing-instances vpls protocols vpls site-range 3
set routing-instances vpls protocols vpls no-tunnel-services
set routing-instances vpls protocols vpls site asia site-identifier 1
set routing-instances vpls protocols vpls site asia interface ge-0/2/2.0
set routing-instances vpls protocols vpls vpls-id 100
set routing-instances vpls protocols vpls bum-hashing
```

Device PE2

```
set interfaces ge-0/0/7 gigether-options 802.3ad ae0
set interfaces ge-0/1/7 gigether-options 802.3ad ae0
set interfaces ge-0/2/3 gigether-options 802.3ad ae0
set interfaces ge-0/2/7 gigether-options 802.3ad ae0
set interfaces ge-2/0/0 gigether-options 802.3ad ae1
set interfaces ge-2/0/1 gigether-options 802.3ad ae1
set interfaces ge-2/0/2 gigether-options 802.3ad ae1
set interfaces ge-2/0/4 encapsulation ethernet-vpls
set interfaces ge-2/0/4 unit 0 family vpls
set interfaces ge-2/0/7 gigether-options 802.3ad ae0
set interfaces ge-2/0/9 unit 0 family inet address 10.10.1.1/30
set interfaces ge-2/0/9 unit 0 family mpls
set interfaces ge-2/1/7 gigether-options 802.3ad ae0
set interfaces ge-2/2/7 gigether-options 802.3ad ae0
set interfaces ge-2/3/7 gigether-options 802.3ad ae0
set interfaces ae0 unit 0 family inet address 10.11.11.2/30
set interfaces ae0 unit 0 family iso
set interfaces ae0 unit 0 family mpls
set interfaces ae1 unit 0 family inet address 10.1.1.1/30
set interfaces ae1 unit 0 family mpls
set protocols rsvp interface all
set protocols rsvp interface fxp0.0 disable
set protocols mpls label-switched-path PE2-to-PE3 from 10.255.19.77
set protocols mpls label-switched-path PE2-to-PE3 to 10.255.19.79
```

```
set protocols mpls label-switched-path PE2-to-PE1 to 10.255.71.214
set protocols mpls interface all
set protocols mpls interface fxp0.0 disable
set protocols bgp group int type internal
set protocols bgp group int local-address 10.255.19.77
set protocols bgp group int family inet any
set protocols bgp group int family l2vpn signaling
set protocols bgp group int neighbor 10.255.71.214
set protocols bgp group int neighbor 10.255.19.79
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-2/0/0.0
set protocols ospf area 0.0.0.0 interface ge-2/0/1.0
set protocols ospf area 0.0.0.0 interface ge-2/0/2.0
set protocols ospf area 0.0.0.0 interface ae0.0
set protocols ospf area 0.0.0.0 interface ae1.0
set protocols ospf area 0.0.0.0 interface all
set protocols ospf area 0.0.0.0 interface fxp0.0 disable
set protocols ldp interface all
set protocols ldp interface fxp0.0 disable
set routing-instances vpls instance-type vpls
set routing-instances vpls interface ge-2/0/4.0
set routing-instances vpls route-distinguisher 65000:1
set routing-instances vpls vrf-target target:65000:1
set routing-instances vpls protocols vpls site-range 3
set routing-instances vpls protocols vpls no-tunnel-services
set routing-instances vpls protocols vpls site 2 site-identifier 2
set routing-instances vpls protocols vpls site 2 interface ge-2/0/4.0
set routing-instances vpls protocols vpls vpls-id 100
set routing-instances vpls protocols vpls bum-hashing
```

Step-by-Step Procedure

To configure Device PE1:

1. Configure Device PE1 interfaces.

```
[edit interfaces]
user@PE1# set ge-0/0/6 gigether-options 802.3ad ae0
user@PE1# set ge-0/1/6 gigether-options 802.3ad ae0
user@PE1# set ge-0/2/2 encapsulation ethernet-vpls
user@PE1# set ge-0/2/2 unit 0 family vpls
```

```
user@PE1# set ge-0/2/3 gigether-options 802.3ad ae0
user@PE1# set ge-0/2/6 gigether-options 802.3ad ae0
user@PE1# set ge-0/3/0 gigether-options 802.3ad ae0
user@PE1# set ge-0/3/1 gigether-options 802.3ad ae0
user@PE1# set ge-0/3/6 gigether-options 802.3ad ae0
user@PE1# set ge-1/0/6 gigether-options 802.3ad ae0
user@PE1# set ge-1/2/6 unit 0 family inet address 10.1.1.2/30
user@PE1# set ae0 unit 0 family inet address 10.11.11.1/30
user@PE1# set ae0 unit 0 family iso
user@PE1# set ae0 unit 0 family mpls
```

2. On Device PE1, configure the packet header data to be used for per-flow load balancing.

```
[edit forwarding-options hash-key family multiservice]
user@PE1# set source-mac
user@PE1# set destination-mac
user@PE1# set payload ip layer-3
```

3. Configure the routing policy on Device PE1.

```
[edit policy-options]
user@PE1# set policy-statement exp-to-fwd term a from community grn-com
user@PE1# set policy-statement exp-to-fwd term a then install-nexthop lsp PE1-to-PE2
user@PE1# set policy-statement exp-to-fwd term a then accept
user@PE1# set policy-options community grn-com members target:65000:1
```

4. Configure Device PE1 routing protocols and MPLS.

```
[edit protocols]
user@PE1# set rsvp interface all
user@PE1# set rsvp interface fxp0.0 disable
user@PE1# set mpls label-switched-path PE1-to-PE2 to 10.255.19.77
user@PE1# set mpls label-switched-path PE1-to-PE3 to 10.255.19.79
user@PE1# set mpls interface all
user@PE1# set mpls interface fxp0.0 disable
user@PE1# set bgp group int type internal
user@PE1# set bgp group int local-address 10.255.71.214
user@PE1# set bgp group int family inet any
user@PE1# set bgp group int family l2vpn signaling
user@PE1# set bgp group int neighbor 10.255.19.77
```

```
user@PE1# set bgp group int neighbor 10.255.19.79
user@PE1# set ospf traffic-engineering
user@PE1# set ospf area 0.0.0.0 interface all
user@PE1# set ospf area 0.0.0.0 interface fxp0.0 disable
```

5. Configure VPLS on Device PE1.

```
[edit routing-instances vpls]
user@PE1# set instance-type vpls
user@PE1# set interface ge-0/2/2.0
user@PE1# set route-distinguisher 65000:1
user@PE1# set vrf-target target:65000:1
user@PE1# set protocols vpls site-range 3
user@PE1# set protocols vpls no-tunnel-services
user@PE1# set protocols vpls site asia site-identifier 1
user@PE1# set protocols vpls site asia interface ge-0/2/2.0
user@PE1# set protocols vpls vpls-id 100
user@PE1# set protocols vpls bum-hashing
```

Step-by-Step Procedure

To configure Device PE2:

1. Configure Device PE2 interfaces.

```
[edit interfaces]
user@PE2# set ge-0/0/7 gigether-options 802.3ad ae0
user@PE2# set ge-0/1/7 gigether-options 802.3ad ae0
user@PE2# set ge-0/2/3 gigether-options 802.3ad ae0
user@PE2# set ge-0/2/7 gigether-options 802.3ad ae0
user@PE2# set ge-2/0/0 gigether-options 802.3ad ae1
user@PE2# set ge-2/0/1 gigether-options 802.3ad ae1
user@PE2# set ge-2/0/2 gigether-options 802.3ad ae1
user@PE2# set ge-2/0/4 encapsulation ethernet-vpls
user@PE2# set ge-2/0/4 unit 0 family vpls
user@PE2# set ge-2/0/7 gigether-options 802.3ad ae0
user@PE2# set ge-2/0/9 unit 0 family inet address 10.10.1.1/30
user@PE2# set ge-2/0/9 unit 0 family mpls
user@PE2# set ge-2/1/7 gigether-options 802.3ad ae0
user@PE2# set ge-2/2/7 gigether-options 802.3ad ae0
user@PE2# set ge-2/3/7 gigether-options 802.3ad ae0
```

```
user@PE2# set ae0 unit 0 family inet address 10.11.11.2/30
user@PE2# set ae0 unit 0 family iso
user@PE2# set ae0 unit 0 family mpls
user@PE2# set ae1 unit 0 family inet address 10.1.1.1/30
user@PE2# set ae1 unit 0 family mpls
```

2. Configure Device PE2 routing protocols and MPLS.

```
[edit protocols]
user@PE2# set rsvp interface all
user@PE2# set rsvp interface fxp0.0 disable
user@PE2# set mpls label-switched-path PE2-to-PE3 from 10.255.19.77
user@PE2# set mpls label-switched-path PE2-to-PE3 to 10.255.19.79
user@PE2# set mpls label-switched-path PE2-to-PE1 to 10.255.71.214
user@PE2# set mpls interface all
user@PE2# set mpls interface fxp0.0 disable
user@PE2# set bgp group int type internal
user@PE2# set bgp group int local-address 10.255.19.77
user@PE2# set bgp group int family inet any
user@PE2# set bgp group int family l2vpn signaling
user@PE2# set bgp group int neighbor 10.255.71.214
user@PE2# set bgp group int neighbor 10.255.19.79
user@PE2# set ospf traffic-engineering
user@PE2# set ospf area 0.0.0.0 interface lo0.0
user@PE2# set ospf area 0.0.0.0 interface ge-2/0/0.0
user@PE2# set ospf area 0.0.0.0 interface ge-2/0/1.0
user@PE2# set ospf area 0.0.0.0 interface ge-2/0/2.0
user@PE2# set ospf area 0.0.0.0 interface ae0.0
user@PE2# set ospf area 0.0.0.0 interface ae1.0
user@PE2# set ospf area 0.0.0.0 interface all
user@PE2# set ospf area 0.0.0.0 interface fxp0.0 disable
user@PE2# set ldp interface all
user@PE2# set ldp interface fxp0.0 disable
```

3. Configure VPLS on Device PE2.

```
[edit routing-instances vpls]
user@PE2# set instance-type vpls
user@PE2# set interface ge-2/0/4.0
user@PE2# set route-distinguisher 65000:1
user@PE2# set vrf-target target:65000:1
```

```
user@PE2# set protocols vpls site-range 3
user@PE2# set protocols vpls no-tunnel-services
user@PE2# set protocols vpls site 2 site-identifier 2
user@PE2# set protocols vpls site 2 interface ge-2/0/4.0
user@PE2# set protocols vpls vpls-id 100
user@PE2# set protocols vpls bum-hashing
```

Results

From configuration mode, confirm your configuration by issuing the `show forwarding-options`, `show interfaces`, `show protocols`, `show policy-options`, and `show routing-options` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Device PE1

```
user@PE1# show forwarding-options
hash-key {
    family multiservice {
        source-mac;
        destination-mac;
        payload {
            ip {
                layer-3;
            }
        }
    }
}
```

```
user@PE1# show interfaces
ge-0/0/6 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/1/6 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/2/2 {
```

```
    encapsulation ethernet-vpls;
    unit 0 {
        family vpls;
    }
}
ge-0/2/3 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/2/6 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/3/0 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/3/1 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/3/6 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-1/0/6 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-1/2/6 {
    unit 0 {
        family inet {
            address 10.1.1.2/30;
        }
    }
}
ae0 {
```

```
    unit 0 {
        family inet {
            address 10.11.11.1/30;
        }
        family iso;
        family mpls;
    }
}
```

```
user@PE1# show protocols
mpls {
    interface all;
    interface fxp0.0 {
        disable;
    }
    label-switched-path PE1-to-PE2 {
        to 10.255.19.77;
    }
    label-switched-path PE1-to-PE3 {
        to 10.255.19.79;
    }
}
rsvp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
bgp {
    group int {
        type internal;
        local-address 10.255.71.214;
        family inet {
            any;
        }
        family l2vpn {
            signaling;
        }
        neighbor 10.255.19.77;
        neighbor 10.255.19.79;
    }
```

```
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}
```

```
user@PE1# show policy-options
policy-statement exp-to-fwd {
    term a {
        from community grn-com;
        then {
            install-nexthop lsp PE1-to-PE2;
            accept;
        }
    }
}
community grn-com members target:65000:1;
```

```
user@PE1# show routing-instances
vpls {
    instance-type vpls;
    interface ge-0/2/2.0;
    route-distinguisher 65000:1;
    vrf-target target:65000:1;
    protocols {
        vpls {
            site-range 3;
            no-tunnel-services;
            site asia {
                site-identifier 1;
                interface ge-0/2/2.0;
            }
            vpls-id 100;
            bum-hashing;
        }
```

```
    }
}
```

Device PE2

```
user@PE2# show interfaces
ge-0/0/7 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/1/7 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/2/3 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-0/2/7 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-2/0/0 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-2/0/1 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-2/0/2 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-2/0/4 {
```

```
    encapsulation ethernet-vpls;
    unit 0 {
        family vpls;
    }
}
ge-2/0/7 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-2/0/9 {
    unit 0 {
        family inet {
            address 10.10.1.1/30;
        }
        family mpls;
    }
}
ge-2/1/7 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-2/2/7 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-2/3/7 {
    gigether-options {
        802.3ad ae0;
    }
}
ae0 {
    unit 0 {
        family inet {
            address 10.11.11.2/30;
        }
        family iso;
        family mpls;
    }
}
ae1 {
```

```
    unit 0 {
        family inet {
            address 10.1.1.1/30;
        }
        family mpls;
    }
}
```

```
user@PE2# show protocols
rsvp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
mpls {
    label-switched-path PE2-to-PE3 {
        from 10.255.19.77;
        to 10.255.19.79;
    }
    label-switched-path PE2-to-PE1 {
        to 10.255.71.214;
    }
    interface all;
    interface fxp0.0 {
        disable;
    }
}
bgp {
    group int {
        type internal;
        local-address 10.255.19.77;
        family inet {
            any;
        }
        family l2vpn {
            signaling;
        }
        neighbor 10.255.71.214;
        neighbor 10.255.19.79;
    }
```

```
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface lo0.0;
        interface ge-2/0/0.0;
        interface ge-2/0/1.0;
        interface ge-2/0/2.0;
        interface ae0.0;
        interface ae1.0;
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}
ldp {
    interface all;
    interface fxp0.0 {
        disable;
    }
}
```

```
user@PE2# show routing-instances
vpls {
    instance-type vpls;
    interface ge-2/0/4.0;
    route-distinguisher 65000:1;
    vrf-target target:65000:1;
    protocols {
        vpls {
            site-range 3;
            no-tunnel-services;
            site 2 {
                site-identifier 2;
                interface ge-2/0/4.0;
            }
            vpls-id 100;
            bum-hashing;
        }
```

```
    }
}
```

Verification

You can monitor the operation of the routing instance by running the `show interfaces ae1.0 extensive` and `monitor interface traffic` commands.

For troubleshooting, you can configure tracing operations for all of the protocols.

RELATED DOCUMENTATION

- *Configuring Point-to-Multipoint LSPs for an MBGP MVPN*
- *Understanding Wildcards to Configure Selective Point-to-Multipoint LSPs for an MBGP MVPN*
- *show interfaces (Aggregated Ethernet)*
- Load Balancing VPLS Non-Unicast Traffic Across Member Links of an Aggregate Interface | **89**

CHAPTER 6

Configuring Other Forwarding Options

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents | 109

Configuring DNS and TFTP Packet Forwarding | 112

Configuring Port-based LAN Broadcast Packet Forwarding | 117

Preventing DHCP Spoofing on MX Series 5G Universal Routing Platforms | 119

Understanding the Hyper Mode Feature on Enhanced MPCs for MX Series Routers and EX9200 Switches | 121

Configuring Hyper Mode on Enhanced MPCs to Speed Up Packet Processing | 124

Unsupported Features and CLI Commands When Hyper Mode Is Enabled | 126

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents

You can configure the router, switch, or interface to act as a Dynamic Host Configuration Protocol (DHCP) and Bootstrap Protocol (BOOTP) relay agent. This means that a locally attached host can issue a DHCP or BOOTP request as a broadcast message. If the router, switch, or interface sees this broadcast message, it relays the message to a specified DHCP or BOOTP server.

You should configure the router, switch, or interface to be a DHCP and BOOTP relay agent if you have locally attached hosts and a distant DHCP or BOOTP server. For MX Series routers connected via IRB, see the note below to prevent BOOTP reply packets from being dropped.

To configure the router or switch to act as a DHCP and BOOTP relay agent, include the `bootp` statement at the `[edit forwarding-options helpers]` hierarchy level:

```
[edit forwarding-options helpers]
bootp {
    client-response-ttl number;
    description text-description;
    interface (interface-name | interface-group) {
        client-response-ttl number;
        description text-description;
        maximum-hop-count number;
        minimum-wait-time seconds;
        no-listen;
        server address {
            logical-system logical-system-name <routing-instance [ <default> routing-instance-
names ]>;
            routing-instance [ <default> routing-instance-names ];
        }
    }
    maximum-hop-count number;
    minimum-wait-time seconds;
    relay-agent-option;
    server server-identifier {
        <logical-system logical-system-name>
<routing-instance [ routing-instance-names ]>;
    }
}
```

To set the description of the BOOTP service, DHCP service, or interface, include the `description` statement.

To set a logical interface or a group of logical interfaces with a specific DHCP relay or BOOTP configuration, include the `interface` statement.

To set the routing instance of the server to forward, include the `routing-instance` statement. You can include as many routing instances as necessary in the same statement.

To stop packets from being forwarded on a logical interface, a group of logical interfaces, or the router or switch, include the `no-listen` statement.

To set the maximum allowed number in the hops field of the BOOTP header, include the `maximum-hop-count` statement. Headers that have a larger number in the hops field are not forwarded. If you omit the `maximum-hop-count` statement, the default value is four hops.

To set the minimum allowed number of seconds in the **secs** field of the BOOTP header, include the `minimum-wait-time` statement. Headers that have a smaller number in the **secs** field are not forwarded. The default value for the minimum wait time is zero (0).

To set the IP address that specify the DHCP or BOOTP server for the router, switch, or interface, include the `server` statement. You can include multiple `server` statements.

To set an IP time-to-live (TTL) value for DHCP response packets sent to a DHCP client, include the `client-response-ttl` statement.

To use the DHCP relay agent option in relayed BOOTP/DHCP messages, include the `relay-agent-option` statement. This option is primarily useful for enabling DHCP forwarding between different VRF routing instances. This option is documented in RFC 3046, *DHCP Relay Agent Information Option*.

You can also configure an individual logical interface to be a DHCP and BOOTP relay agent if you have locally attached hosts and a remote DHCP or BOOTP server connected to one of the router's or switch's interfaces. For more information, see the Junos OS Administration Library for Routing Devices.

The following example demonstrates a BOOTP relay agent configuration.

```
user@host# show forwarding-options
helpers {
    bootp {
        description "dhcp relay agent global parameters";
        server 192.168.55.44;
        server 172.16.0.3 routing-instance c3;
        maximum-hop-count 10;
        minimum-wait-time 8;
        interface {
            fe-1/3/0 {
```

```
                description "use this info for this interface";
                server 10.10.10.10;
                server 192.168.14.14;
                maximum-hop-count 11;
                minimum-wait-time 3;
            }
            fe-1/3/1 {
                no-listen;  ###ignore DHCPDISCOVER messages on this interface
            }
            all {
                description "globals apply to all other interfaces";
            }
        }
    }
}
```

BEST PRACTICE: To use `bootp helper` on a MX Series router (MX80, MX240, MX480 and MX960) connected via IRB, you may need to take steps to ensure that DHCP discover packets (the bootp reply) are sent to clients and received as expected. Otherwise, **bootp** replies may be dropped because the DHCP client is clearing the broadcast bit in the discover packet, or because the DHCP server is stripping **option-82** flags from the **offer**.

This happens when the IRB interface is a layer 3 (logical) interface associated with a bridge domain that has multiple layer 2 (physical) interfaces associated with it. In such cases, if the **offer** from the DHCP server is unicast and doesn't include an ingress interface identifying the physical interface on which the **discovery** packet was received, the MX router won't be able to determine an interface for sending out **offers**.

1. Enable `broadcast` on the IRB interface to flood **discovery** frames from all physical interfaces in the bridge domain. For example,

```
user@host# edit forwarding-options helpers bootp interface irb.o
    broadcast;
        server 202.67.4.1;
    }
```

or,

2. Enable `relay-agent-option` on the bootp helper. For example,

```
user@host# edit forwarding-options helpers bootp
    relay-agent-option;
        server 202.67.4.1;
    }
```

3. Configure the IRB interface connected to the DHCP server so it echoes **option-82** flags back to the router. This will ensure that the **option-82** string, which identifies the interface used by the router, is preserved.

Configuring DNS and TFTP Packet Forwarding

IN THIS SECTION

- Tracing BOOTP, DNS, and TFTP Forwarding Operations | **113**
- Example: Configuring DNS Packet Forwarding | **116**

You can configure the router or switch to support Domain Name System (DNS) and Trivial File Transfer Protocol (TFTP) packet forwarding for IPv4 traffic, which allows clients to send DNS or TFTP requests to the router or switch. The responding DNS or TFTP server recognizes the client address and sends a response directly to that address. By default, the router or switch ignores DNS and TFTP request packets.

To enable DNS or TFTP packet forwarding, include the `helpers` statement at the `[edit forwarding-options]` hierarchy level:

```
[edit forwarding-options]
helpers {
    domain {
        description text-description;
        interface interface-name {
            description text-description;
            no-listen;
            server [ addresses {
```

```
                logical-system logical-system-name;
                routing-instance instance-name;
            }
        }
    }
    tftp {
        description text-description;
        interface interface-name {
            description text-description;
            no-listen;
            server address;
            server logical-system name < [ routing-instance routing-instance-names ] >;
            server < [ routing-instance routing-instance-names ] >;
        }
    }
}
```

To set domain packet forwarding, include the `domain` statement.

To set the description of the DNS or TFTP service, include the `description` statement.

To set TFTP packet forwarding, include the `tftp` statement.

To set a DNS or TFTP server (with an IPv4 address), include the `server` statement. Use one address for either a global configuration or for each interface.

To set the routing instance of the server to forward, include the `routing-instance` statement. You can include as many routing instances as necessary in the same statement.

To disable recognition of DNS or TFTP requests on one or more interfaces, include the `no-listen` statement. If you do not specify at least one interface with this statement, the forwarding service is global to all interfaces on the router or switch.

The following sections discuss the following:

Tracing BOOTP, DNS, and TFTP Forwarding Operations

BOOTP, DNS, and TFTP forwarding tracing operations track all BOOTP, DNS, and TFTP operations and record them in a log file. The logged error descriptions provide detailed information to help you solve problems faster.

By default, nothing is traced. If you include the `traceoptions` statement at the `[edit forwarding-options helpers]` hierarchy level, the default tracing behavior is the following:

- Important events are logged in a file called **fud** located in the **/var/log** directory.
- When the file **fud** reaches 128 kilobytes (KB), it is renamed **fud.0**, then **fud.1**, and so on, until there are 3 trace files. Then the oldest trace file (**fud.2**) is overwritten. (For more information about how log files are created, see the System Log Explorer.)
- Log files can be accessed only by the user who configures the tracing operation.

You cannot change the directory (**/var/log**) in which trace files are located. However, you can customize the other trace file settings by including the following statements at the `[edit forwarding-options helpers]` hierarchy level:

```
[edit forwarding-options helpers]
traceoptions {
    file filename <files number> <match regular-expression> <size size> <world-readable | no-
world-readable>;
    flag {
        address;
        all;
        config;
        domain;
        ifdb;
        io;
        main;
        port;
        rtsock;
        tftp;
        trace;
        ui;
        util;
    }
    level severity-level;
    no-remote-trace;
}
```

These statements are described in the following sections:

Configuring the Log Filename

By default, the name of the file that records trace output is **fud**. You can specify a different name by including the `file` *filename* statement at the `[edit forwarding-options helpers traceoptions]` hierarchy level:

```
[edit forwarding-options helpers traceoptions]
file filename;
```

Configuring the Number and Size of Log Files

By default, when the trace file reaches 128 kilobytes (KB) in size, it is renamed ***filename*.0**, then ***filename*.1**, and so on, until there are three trace files. Then the oldest trace file (***filename*.2**) is overwritten.

You can configure the limits on the number and size of trace files by including the following statements at the `[edit forwarding-options helpers traceoptions]` hierarchy level:

```
[edit forwarding-options helpers traceoptions]
file files number size size;
```

For example, set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation (***filename***) reaches 2 MB, ***filename*** is renamed ***filename*.0**, and a new file called ***filename*** is created. When the new ***filename*** reaches 2 MB, ***filename*.0** is renamed ***filename*.1** and ***filename*** is renamed ***filename*.0**. This process repeats until there are 20 trace files. Then the oldest file (***filename*.19**) is overwritten by the newest file (***filename*.0**).

The number of files can be from 2 through 1000 files. The file size of each file can be from 10 KB through 1 gigabyte (GB).

Configuring Access to the Log File

By default, log files can be accessed only by the user who configures the tracing operation.

To specify that any user can read all log files, include the **world-readable** option with the `file` statement at the `[edit forwarding-options helpers traceoptions]` hierarchy level:

```
[edit forwarding-options helpers traceoptions]
file world-readable;
```

To explicitly set the default behavior, include the **no-world-readable** option with the `file` statement at the `[edit forwarding-options helpers traceoptions]` hierarchy level:

```
[edit forwarding-options helpers traceoptions]
file no-world-readable;
```

Configuring a Regular Expression for Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged events.

You can refine the output by including the **match** option with the `file` statement at the `[edit forwarding-options helpers traceoptions]` hierarchy level and specifying a regular expression (regex) to be matched:

```
[edit forwarding-options helpers traceoptions]
file filename match regular-expression;
```

Example: Configuring DNS Packet Forwarding

Enable DNS packet request forwarding to all interfaces on a router except **t1-1/1/2** and **t1-1/1/3**:

```
[edit forwarding-options helpers]
dns {
    server 10.10.10.30;
    interface {
        t1-1/1/2 {
            no-listen;
            server 10.10.10.9;
        }
        t1-1/1/3 {
            no-listen;
            server 10.10.10.4;
        }
    }
}
```

Configuring Port-based LAN Broadcast Packet Forwarding

You can enable a router or switch to forward LAN broadcast traffic on custom UDP ports to specified servers by configuring *port helpers* with the `[edit forwarding-options helpers] port` configuration statement. Port helpers are also referred to as port forwarding or UDP broadcast packet forwarding services. When you configure a port helper, the router or switch listens for incoming UDP traffic for the configured port with destination Layer 2 MAC and Layer 3 IP broadcast addresses, and forwards the packets as unicast traffic to a configured server.

Port helpers forward the traffic for configured ports transparently, without considering the application layer protocols in the packets being forwarded. However, you cannot configure a port helper to forward traffic for standard ports used by services such as BOOTP, DNS and TFTP. These services have their own explicit packet forwarding helper configuration options (see "helpers" on page 212 and "Configuring DNS and TFTP Packet Forwarding" on page 112).

You can configure port helpers to listen for and forward broadcast traffic for a configured port using any of the following scopes:

- Global scope—Forward incoming broadcast traffic on the port to a configured destination server IP address.

  Configure a global port helper using only the `server` configuration option, without specifying a particular interface. The port helper listens for incoming traffic on any interfaces to forward to the configured server. For example:

  ```
  set forwarding-options helpers port 1300 server 10.20.30.40
  ```

- VLAN-specific scope—Forward incoming broadcast traffic on the port from a configured VLAN to a configured destination server IP address.

  Configure a VLAN-specific port helper using the `interface` statement with an IRB interface name for a VLAN, and the `server` statement. The port helper listens for incoming traffic from interfaces in the VLAN to forward to the configured server. For example:

  ```
  set forwarding-options helpers port 1064 interface irb.100 server 192.0.2.50
  ```

- Interface-specific scope—Forward incoming broadcast traffic on the port from a configured Layer 3 interface to a configured destination server IP address.

Configure an interface-specific port helper using the `interface` statement with a Layer 3 interface name, and the `server` statement. The port helper listens for incoming traffic only from the configured interface to forward to the configured server. For example:

```
set forwarding-options helpers port 1064 interface ge-0/0/3 server 192.0.2.50
```

For any scope, optionally use the `description` statement to label or describe the configured forwarding service.

In releases prior to Junos OS Release 17.2, you can configure only one destination server for a given port number. Starting in Junos OS Release 17.2R1, you can configure forwarding traffic to multiple servers for a given port in any port helper scope. To configure forwarding the traffic on a specified port to multiple destination servers, include multiple configuration items for the port and each server (or interface and server). For example, in the global scope:

```
set forwarding-options helpers port 1300 server 10.20.30.4
set forwarding-options helpers port 1300 server 10.20.30.5
set forwarding-options helpers port 1300 server 10.20.30.6
```

To temporarily disable listening on a configured port from a configured interface, include the `no-listen` option with the configured item, as follows:

```
set forwarding-options helpers port port-number interface interface-name server address no-
listen
```

To remove a configured port helper service from a router or switch, delete the configured `port` number item, as follows:

```
delete forwarding-options helpers port <port-number>
```

If multiple servers are configured for a particular port, to remove any or all such forwarding services, you must delete each configured port and server item individually. For example:

```
delete forwarding-options helpers port 1300 server 10.20.30.4
delete forwarding-options helpers port 1300 server 10.20.30.5
delete forwarding-options helpers port 1300 server 10.20.30.6
```

Release History Table

Release	Description
17.2R1	Starting in Junos OS Release 17.2R1, you can configure forwarding traffic to multiple servers for a given port in any port helper scope.

RELATED DOCUMENTATION

port (Packet Forwarding) | 288

server (DNS, Port, and TFTP Service) | 310

interface (DNS, Port, and TFTP Packet Forwarding or Relay Agent) | 235

Preventing DHCP Spoofing on MX Series 5G Universal Routing Platforms

A problem that sometimes occurs with DHCP is *DHCP spoofing*. in which an untrusted client floods a network with DHCP messages. Often these attacks utilize source IP address spoofing to conceal the true source of the attack.

DHCP snooping helps prevent DHCP spoofing by copying DHCP messages to the control plane and using the information in the packets to create anti-spoofing filters. The anti-spoofing filters bind a client's MAC address to its DHCP-assigned IP address and use this information to filter spoofed DHCP messages. In a typical topology, a carrier edge router (in this function also referred to as the broadband services router [BSR]) connects the DHCP server and the MX Series router (or broadband services aggregator [BSA]) performing the snooping. The MX Series router connects to the client and the BSR.

DHCP snooping works as follows in the network topology mentioned above:

1. The client sends a DHCP discover message to obtain an IP address from the DHCP server.
2. The BSA intercepts the message and might add option 82 information specifying the slot, port, VPI/VCI, and so on.
3. The BSA then sends the DHCP discover message to the BSR, which converts it to a unicast packet and sends it to the DHCP server.
4. The DHCP server looks up the client's MAC address and option 82 information in its database. A valid client is assigned an IP address, which is returned to the client using a DHCP offer message. Both the BSR and BSA send this message upstream to the client.

5. The client examines the DHCP offer, and if it is acceptable, issues a DHCP request message that is sent to the DHCP server through the BSA and BSR.
6. The DHCP server confirms that the IP address is still available. If it is, the DHCP server updates its local tables and sends a DHCP ACK message to the client.
7. The BSR receives the DHCP ACK message and passes the message to the BSA.
8. The BSA creates an anti-spoofing filter by binding the IP address in the ACK message to the MAC address of the client. After this point, any DHCP messages from this IP address that are not bound to the client's MAC address are dropped.
9. The BSA sends the ACK message to the client so that the process of assigning a IP address can be completed.

You configure DHCP snooping by including within a DHCP group the appropriate interfaces of the BSA:

```
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name forwarding-options
dhcp-relay group group-name]
interface interface-name;
```

In a VPLS environment, DHCP requests are forwarded over pseudowires. You can configure DHCP snooping over VPLS at the `[edit routing-instances routing-instance-name]` hierarchy level.

DHCP snooping works on a per learning bridge basis in bridge domains. Each learning domain must have an upstream interface configured. This interface acts as the flood port for DHCP requests coming from the client side. DHCP requests are be forwarded across learning domains in a bridge domain. You can configure DHCP snooping on bridge domains at the `[edit routing-instances routing-instance-name bridge-domains bridge-domain-name]` hierarchy level.

RELATED DOCUMENTATION

Preventing DHCP Spoofing

Understanding the Hyper Mode Feature on Enhanced MPCs for MX Series Routers and EX9200 Switches

Starting with Junos OS Release 15.1, enhanced MPCs can be configured to support increased packet processing rates. Enhanced MPCs include these models: MPC3E, MPC4E, MPC5E, MPC6E, MPC7E-MRATE, MPC7E-10G, MX2K-MPC8E, and MX2K-MPC9E.

Starting with Junos OS Release 18.2R1, MPC JNP10K-LC2101 can be configured to support increased packet processing rates. A higher rate of processing of data packets results in the optimization of the lifetime of a data packet. Optimization of the data packet lifetime enables the network device (a router or a switch) to provide better performance and throughput.

To enable the device to support increased packet processing rates, you must configure the hyper mode feature. After configuring the hyper mode feature, you must reboot the device for the changes to take effect.

When you configure the hyper mode feature on the device, the configured mode changes from normal mode to hyper mode. However, because the configuration does not take effect until you reboot the device the current mode of the device remains as normal mode. The current mode changes from normal mode to hyper mode after you reboot the device. If the hyper mode feature is not configured, thedevice processes data packets in normal mode.

NOTE:

- You can enable the hyper mode feature only if the network-service mode on the device is configured as either `enhanced-ip` or `enhanced-ethernet`.
- Hyper-mode is the default forwarding mode on the SCBE3-MX. If your deployment does not need hyper-mode, disable hyper-mode using the `set forwarding-options no-hyper-mode` cli command before installing the Routing Engine into the SCBE3-MX. See also: SCBE3-MX Description

Table 2 on page 122 displays the values of the current and configured mode based on the hyper mode configuration and system reboot.

Table 2: Current Mode and Configured Mode Values Based on Hyper mode Configuration

Action	Current Mode	Configured Mode
Hyper mode is configured but the device is not rebooted.	Normal mode	Hyper mode
Hyper mode is configured and device is rebooted.	Hyper mode	Hyper mode
Hyper mode configuration is removed and device is not rebooted.	Hyper mode	Normal mode
Hyper mode configuration is removed and device is rebooted.	Normal mode	Normal mode

When you configure hyper mode, the following features are not supported:

- Creation of Virtual Chassis
- Forwarding class accounting (enhanced mode)
- Interoperability with legacy DPCs, including MS-DPCs. The MPC in hyper mode accepts and transmits data packets only from other existing MPCs.
- Interoperability with non-Ethernet MICs and non-Ethernet Interfaces such as channelized interfaces, multilink interfaces, and SONET interfaces.
- Junos Fusion
- Junos Node Slicing
- Padding of Ethernet frames with VLAN is not supported in Junos OS releases prior to 19.2R1.
- Precision Time Protocol
- Provider Backbone Bridging (PBB) and Ethernet VPN (EVPN)
- Sending Internet Control Message Protocol (ICMP) redirect messages is not supported in Junos OS releases prior to 19.2R1. In versions prior to 19.2R1, ICMP redirects are disabled by default and cannot be re-enabled in hyper mode.
- Termination or tunneling of all subscriber-based services.
- Collecting interface family statistics for IPv4 and IPv6 is not supported in Junos OS releases prior to 19.2R1. The interface family statistics are collected by using the `show interfaces statistics detail` *`interface-name`* command.

After you configure the hyper mode feature and reboot the device, existing MPCs that do not support the hyper mode feature, such as MPC1, MPC2, and MPC3, power on in normal mode. Also, when you have installed MICs and PICs on MPCs that are in normal mode when the hyper mode feature is enabled, those MICs and PICs do not power on. Following is a list of the MICs and PICs that do not power on:

- Channelized E1/T1 Circuit Emulation MIC
- Channelized E1/T1 Circuit Emulation MIC (H)
- *Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP*
- Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (H)
- Channelized SONET/SDH OC3/STM1 (Multi-Rate) MICs with SFP
- DS3/E3 MIC
- SONET/SDH OC3/STM1 (Multi-Rate) MICs with SFP
- SONET/SDH OC192/STM64 MIC with XFP
- Channelized OC48/STM16 Enhanced IQ (IQE) PIC with SFP

Release History Table

Release	Description
18.2R1	Starting with Junos OS Release 18.2R1, MPC JNP10K-LC2101 can be configured to support increased packet processing rates.
15.1	Starting with Junos OS Release 15.1, enhanced MPCs can be configured to support increased packet processing rates.

RELATED DOCUMENTATION

Configuring Hyper Mode on Enhanced MPCs to Speed Up Packet Processing | **124**

Unsupported Features and CLI Commands When Hyper Mode Is Enabled | **126**

show forwarding-options hyper-mode | **355**

hyper-mode (forwarding-options) | **219**

Configuring Hyper Mode on Enhanced MPCs to Speed Up Packet Processing

Starting with Junos OS Release 15.1, enhanced MPCs can be configured to support increased packet processing rates. Enhanced MPCs include these models: MPC3E, MPC4E, MPC5E, MPC6E,MPC7E-MRATE, MPC7E-10G, MX2K-MPC8E, and MX2K-MPC9E.

Starting with Junos OS Release 18.2R1, JNP10K-LC2101 MPC can be configured to support increased packet processing rates.A higher rate of processing of data packets results in the optimization of the lifetime of a data packet. Optimization of the data packet lifetime enables the network device (a router or a switch) to provide better performance and throughput.

To configure the device to support increased packet processing rates, you must configure the hyper mode feature. After configuring the hyper mode feature, you must reboot the device for the changes to take effect. If the hyper mode feature is not configured, the device processes data packets in normal mode.

NOTE: You can enable the hyper mode feature only if the network-service mode on the device is configured as either `enhanced-ip` or `enhanced-ethernet`.

To configure hyper mode on enhanced MPCs to speed up packet processing:

1. Configure hyper mode by including the `forwarding-options hyper-mode` statement at the [edit] hierarchy level.

```
[edit]
user@host# set forwarding-options hyper-mode
```

2. After configuring hyper mode, commit the configuration.

```
[edit]
user@host# commit
```

NOTE: After configuring hyper mode and committing the configuration, the configured mode changes to `hyper-mode` but the current mode remains as `normal mode`. The device displays the following warning message after you commit the configuration:

```
[edit forwarding-options] 'hyper-mode' WARNING: forwarding-options hyper-mode configuration changed. A system reboot is mandatory. Please reboot the system NOW. Continuing without a reboot might result in unexpected system behavior. commit complete
```

When hyper-mode configuration is enabled on MX960 devices and reboot is performed, the backup routing engine might not come back up in correct mode.

To recover the backup RE, perform `commit force`, to get the same configuration as the master routing engine. However, to get the enhanced IP configuration on the backup routing engine for the network services, after syncing the script, perform reboot using the command `request routing-engine login other-routing-engine`. Now both the routing engines will come up in enhanced IP mode.

3. Reboot the device for the configuration to take effect.

```
user@host> request system reboot
```

Release History Table

Release	Description
18.2R1	Starting with Junos OS Release 18.2R1, JNP10K-LC2101 MPC can be configured to support increased packet processing rates.
15.1	Starting with Junos OS Release 15.1, enhanced MPCs can be configured to support increased packet processing rates.

RELATED DOCUMENTATION

Understanding the Hyper Mode Feature on Enhanced MPCs for MX Series Routers and EX9200 Switches | **121**

Unsupported Features and CLI Commands When Hyper Mode Is Enabled | **126**

show forwarding-options hyper-mode | **355**

hyper-mode (forwarding-options) | **219**

Unsupported Features and CLI Commands When Hyper Mode Is Enabled

Table 3 on page 126 lists the features and corresponding CLI commands that are not supported when the hyper mode feature is enabled. Also, the table lists the error messages displayed when you use the unsupported commands.

Table 3: Unsupported Features and CLI Commands When Hyper Mode Is Enabled

Features	Commands	Error Message
Virtual Chassis	`set virtual-chassis preprovisioned`	`To configure virtual-chassis, 'forwarding-options hyper-mode' should not be configured`
	`set virtual-chassis member` *`member-id`* `role` *`role`* `serial-number` *`ser_num`*	
	`set virtual-chassis no-split-detection`	
ICMP Redirects Supported in Junos OS Release 19.2R1 and later releases.	`set system no-redirects`	Supported in Junos OS Release 19.2R1 and later releases. In Junos OS versions prior to 19.2R1, the following error message was displayed: `To configure system no-redirects, 'forwarding-options hyper-mode' should not be configured`
	`set system no-redirects-ipv6`	Supported in Junos OS Release 19.2R1 and later releases. In Junos OS versions prior to 19.2R1, the following error message was displayed: `To configure system no-redirects-ipv6, 'forwarding-options hyper-mode' should not be configured`

Table 3: Unsupported Features and CLI Commands When Hyper Mode Is Enabled *(Continued)*

Features	Commands	Error Message
	`set interfaces *interface-name* unit *logical-unit-number* family *family* no-redirects`	Supported in Junos OS Release 19.2R1 and later releases. In Junos OS versions prior to 19.2R1, the following error message was displayed: `To configure family inet no-redirects, 'forwarding-options hyper-mode' should not be configured`
PPPoE	`set protocols pppoe service-name-tables *table-name*`	`To configure pppoe, 'forwarding-options hyper-mode' should not be configured`
	`set dynamic-profiles *profile-name* interfaces *interface-name* unit *unit* family pppoe`	`To configure family pppoe, 'forwarding-options hyper-mode' should not be configured`
	`set dynamic-profiles *profile-name* interfaces demux0 unit *unit* family pppoe`	`To configure family pppoe, 'forwarding-options hyper-mode' should not be configured`
L2TP	`set access tunnel-profile *profile-name* tunnel *tunnel-id* tunnel-type l2tp`	`To configure l2tp, 'forwarding-options hyper-mode' should not be configured`
	`set services l2tp`	`To configure services l2tp, 'forwarding-options hyper-mode' should not be configured`
Forwarding class accounting (enhanced mode)	For more information and CLIs, see CoS-Based Interface Counters for IPv4 or IPv6 Aggregate on Layer 2 and forwarding-class-accounting.	--
Junos Node Slicing	For CLIs, see Setting Up Junos Node Slicing.	--

Table 3: Unsupported Features and CLI Commands When Hyper Mode Is Enabled *(Continued)*

Features	Commands	Error Message
Junos Fusion	For CLIs, see Junos Fusion Provider Edge User Guide and Junos Fusion Enterprise User Guide.	--
Provider Backbone Bridging (PBB) and Ethernet VPN (EVPN)	For CLIs, see EVPN User Guide.	--
Precision Time Protocol	For more information, see Configuring Precision Time Protocol	—

RELATED DOCUMENTATION

- Configuring Hyper Mode on Enhanced MPCs to Speed Up Packet Processing | **124**
- Understanding the Hyper Mode Feature on Enhanced MPCs for MX Series Routers and EX9200 Switches | **121**
- show forwarding-options hyper-mode | **355**
- hyper-mode (forwarding-options) | **219**

CHAPTER 7

Configuration Statements

accounting | 133

aggregation | 135

autonomous-system-type | 137

bootp | 139

bum-hashing | 141

cflowd (Discard Accounting) | 142

cflowd (Flow Monitoring) | 144

client-address | 146

client-response-ttl | 148

description (Forwarding Options) | 150

dhcp-relay (DHCP Spoofing Prevention) | 151

disable (Forwarding Options) | 153

domain | 155

ecmp-local-bias | 156

enhanced-hash-key | 158

export-format | 164

family (Filtering) | 166

family (Monitoring) | 168

family (Port Mirroring) | 169

family (Sampling) | 172

family inet | 174

family mpls | 177

family multiservice | 181

file (Extended DHCP Relay Agent and Helpers Trace Options) | 184

file (Sampling) | 186

file (Trace Options) | 187

filename (Sampling) | 189

files (Sampling and Traceoptions) | 190

filter (IPv4, IPv6, and MPLS) | 192

filter (VPLS) | 194

flood | 195

flow-active-timeout | 197

flow-export-destination | 199

flow-inactive-timeout | 200

flow-server | 202

group (DHCP Spoofing Prevention) | 204

gtp-tunnel-endpoint-identifier | 206

hash-key (Forwarding Options) | 208

helpers | 212

hosted-service-identifier | 216

hosted-services | 218

hyper-mode (forwarding-options) | 219

indexed-load-balance | 221

input (Forwarding Table) | 223

input (Port Mirroring) | 225

input (Sampling) | 226

instance | 228

interface (Accounting or Sampling) | 230

interface (BOOTP) | 232

interface (DHCP Spoofing Prevention) | 234

interface (DNS, Port, and TFTP Packet Forwarding or Relay Agent) | 235

interface (Monitoring) | 237

interface (Next-Hop Group) | 239

interface (Port Mirroring) | 241

l2tp-tunnel-session-identifier | 243

link-layer-broadcast-inet-check | 244

load-balance (Forwarding Options) | 246

load-balance-group | 249

local-bias | 250

local-dump | 252

max-packets-per-second | 253

maximum-hop-count | 255

maximum-packet-length | 257

minimum-wait-time | 259

mirror-once | 261

monitoring | 262

next-hop (Forwarding Options) | 264

next-hop-group (Forwarding Options) | 266

next-hop-group | 268

no-filter-check | 270

no-listen | 272

output (Accounting) | 273

output (Forwarding Table) | 275

output (Monitoring) | 277

output (Port Mirroring) | 278

output (Sampling) | 281

per-flow | 283

per-prefix | 285

port (cflowd) | 286

port (Packet Forwarding) | 288

port-mirroring | 291

rate (Forwarding Options) | 296

relay-agent-option | 298

route-accounting | 300

run-length | 302

sampling (Forwarding Options) | 304

server (DHCP and BOOTP Relay Agent) | 308

server (DNS, Port, and TFTP Service) | 310

server-address (Hosted Services) | 312

server-profile | 313

server-profile (Active Flow Monitoring) | 315

size (Sampling and Traceoptions) | 316

source-checking | 318

stamp | 320

tftp | 321

traceoptions (DNS, Port, and TFTP Packet Forwarding) | 323

traceoptions (Port Mirroring and Traffic Sampling) | 326

vector routing | 328

vector-routing | 328

accounting

IN THIS SECTION

- Syntax | 133
- Hierarchy Level | 134
- Description | 134
- Required Privilege Level | 134
- Release Information | 134

Syntax

```
accounting group-name {
    output {
        aggregate-export-interval seconds;
        cflowd [ hostnames ] {
            aggregation {
                autonomous-system;
                destination-prefix;
                protocol-port;
                source-destination-prefix {
                    caida-compliant;
                }
                source-prefix;
            }
            autonomous-system-type (origin | peer);
            port port-number;
            version format;
        }
        flow-active-timeout seconds;
        flow-inactive-timeout seconds;
        interface interface-name {
            engine-id number;
            engine-type number;
            source-address address;
```

```
        }
    }
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Specify discard accounting instance name and options.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Discard Accounting | 31

aggregation

IN THIS SECTION

- Syntax | 135
- Hierarchy Level | 135
- Description | 136
- Options | 136
- Required Privilege Level | 136
- Release Information | 136

Syntax

```
aggregation {
    autonomous-system;
    destination-prefix;
    protocol-port;
    source-destination-prefix {
        caida-compliant;
    }
    source-prefix;
}
```

Hierarchy Level

```
[edit forwarding-options accounting output hostname],
[edit forwarding-options sampling family (inet | inet6 | mpls)output flow-server hostname]
```

Description

For cflowd version 8 only, specify the type of data to be aggregated; cflowd records and sends only those flows that match the specified criteria.

Options

autonomous-system—Aggregate by autonomous system (AS) number.

caida-compliant—Record source and destination mask length values in compliance with the Version 2.1b1 release of the cflowd application from the Cooperative Association for Internet Data Analysis (CAIDA). If this statement is not configured, Junos OS records source and destination mask length values in compliance with the *cflowd Configuration Guide*, dated August 30, 1999.

destination-prefix—Aggregate by destination prefix.

protocol-port—Aggregate by protocol and port number.

source-destination-prefix—Aggregate by source and destination prefix.

source-prefix—Aggregate by source prefix.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Directing Traffic Sampling Output to a Server Running the cflowd Application | 12

autonomous-system-type

IN THIS SECTION

- Syntax | 137
- Hierarchy Level | 137
- Description | 137
- Default | 138
- Options | 138
- Required Privilege Level | 138
- Release Information | 138

Syntax

```
autonomous-system-type (origin | peer);
```

Hierarchy Level

```
[edit forwarding-options accountingoutput cflowd hostname],
[edit forwarding-options sampling family (inet | inet6 | mpls)output flow-server hostname]
```

Description

Specify the type of AS numbers that cflowd exports.

Default

origin

Options

origin—Export origin AS numbers of the packet source address in the Source Autonomous System cflowd field.

peer—Export peer AS numbers through which the packet passed in the Source Autonomous System cflowd field.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Directing Traffic Sampling Output to a Server Running the cflowd Application | 12

bootp

IN THIS SECTION

- Syntax | 139
- Hierarchy Level | 140
- Description | 140
- Options | 140
- Required Privilege Level | 140
- Release Information | 140

Syntax

```
bootp {
    client-response-ttl number;
    description text-description;
    interface (interface-name | interface-group) {
        client-response-ttl number;
        description text-description;
        maximum-hop-count number;
        minimum-wait-time seconds;
        no-listen;
        server address {
            logical-system logical-system-name <routing-instance [ <default> routing-instance-
names ]>;
            routing-instance [ <default> routing-instance-names ];
        }
    }
    maximum-hop-count number;
    minimum-wait-time seconds;
    relay-agent-option;
    server address {
        <logical-system logical-system-name> <routing-instance
 [ routing-instance-names ]>;
```

```
    }
}
```

Hierarchy Level

```
[edit forwarding-options helpers]
```

Description

Configures a router, switch, or interface to act as a Dynamic Host Configuration Protocol (DHCP) or bootstrap protocol (BOOTP) relay agent. For MX Series (MX80, MX240, MX480 and MX960) routers connected via IRB, see "Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents" on page 109 for information on how to prevent BOOTP reply packets from being dropped.

DHCP relaying is disabled.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents | 109

bum-hashing

IN THIS SECTION

- Syntax | 141
- Hierarchy Level | 141
- Description | 141
- Required Privilege Level | 142
- Release Information | 142

Syntax

```
bum-hashing;
```

Hierarchy Level

```
[edit logical-systems logical-system-name routing-instances routing-instance-name protocols
vpls],
[edit routing-instances routing-instance-name protocols vpls]
```

Description

Load balance VPLS BUM (broadcast, unknown, and multicast) traffic across all members of an aggregate interface for the routing instance.

NOTE: On MX Series routers, the bum-hashing configuration is not required when the router is operating in `enhanced-ip` or `enhanced-ethernet` modes. If set, the configuration is ignored.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Load Balancing VPLS Non-Unicast Traffic Across Member Links of an Aggregate Interface | 89

Example: Configuring Multicast Load Balancing over Aggregated Ethernet Links | 90

cflowd (Discard Accounting)

IN THIS SECTION

- Syntax | 143
- Hierarchy Level | 143
- Description | 143
- Options | 143
- Required Privilege Level | 144
- Release Information | 144

Syntax

```
cflowd hostname {
    aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
            caida-compliant;
        }
        source-prefix;
    }
    autonomous-system-type (origin | peer);
    port port-number;
    version format;
}
```

Hierarchy Level

```
[edit forwarding-options accounting group-name output]
```

Description

Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility cfdcollect.

You can configure up to one version 5 and one version 8 flow format at the `[edit forwarding-options accounting group-name output]` hierarchy level.

Options

hostname—The IP address or identifier of the host system (the workstation running the cflowd utility).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Directing Traffic Sampling Output to a Server Running the cflowd Application | 12

cflowd (Flow Monitoring)

IN THIS SECTION

- Syntax | 145
- Hierarchy Level | 145
- Description | 145
- Options | 146
- Required Privilege Level | 146
- Release Information | 146

Syntax

```
cflowd hostname {
    port port-number;
}
```

Hierarchy Level

```
[edit forwarding-options monitoring group-namefamily inet output]
```

Description

Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility cfdcollect.

You can configure up to eight version 5 flow formats at the `[edit forwarding-options monitoring group-name output]` hierarchy level. Version 8 flow formats are not supported for flow-monitoring applications.

J-Flow does not require a license on SRX devices. J-Flow versions 5, 8, and 9 are supported on SRX series devices.

J-Flow version 9 on standalone devices is supported as of:

- SRX Branch devices (SRX1x0, SRX2x0, SRX550, SRX650), Junos 10.4
- SRX-HE devices (SRX1400, SRX3x00, SRX5x00), Junos 12.1X45-D10
- SRX3x0 & SRX550M, Junos 15.1X49-D30
- SRX1500, SRX4100, SRX4200, vSRX, 15.1X49-D80
- SRX4600, Junos 17.4R1-S1

J-Flow version 9 on chassis cluster devices as of:

- SRX Branch devices (SRX-300/320/340/345/380/550HM), Junos 20.1R1
- SRX-HE devices (SRX1400, SRX3x00, SRX5x00), Junos 12.1X45-D10

- SRX1500, SRX4100, SRX4200, vSRX, Junos 15.1X49-D80
- SRX4600, Junos 17.4R1-S1

Options

hostname—The IP address or identifier of the host system (the workstation running the cflowd utility).

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Directing Traffic Sampling Output to a Server Running the cflowd Application | 12

Configuring Passive Flow Monitoring | 35

client-address

IN THIS SECTION

Syntax | 147

Hierarchy Level | 147

- Description | 147
- Options | 147
- Required Privilege Level | 148
- Release Information | 148

Syntax

```
client-address ipv4-address;
```

Hierarchy Level

```
[edit services hosted-services server-profile server-profile-name]
```

Description

Configure the source address to include in the header of each sampled packet. You must specify an IPv4 address. You can also specify the loopback address or the management interface address as the client address.

Options

ipv4-address—IPv4 address of the client.

- **Default:** 0.0.0.0

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

client-response-ttl

IN THIS SECTION

- Syntax | 148
- Hierarchy Level | 149
- Description | 149
- Options | 149
- Required Privilege Level | 149
- Release Information | 149

Syntax

```
client-response-ttl number;
```

Hierarchy Level

```
[edit forwarding-options helpers bootp],
[edit forwarding-options helpers bootp interface (interface-name | interface-group)]
```

Description

Set the IP time-to-live (TTL) value in BOOTP response messages sent to a BOOTP client. If you do not include the `client-response-ttl` statement, the default is to leave the TTL field unchanged.

Options

number—IP time-to-live (TTL) value.

- **Range:** 1 through 255
- **Default:** Leave the TTL field in the BOOTP response message unchanged.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents | 109

description (Forwarding Options)

IN THIS SECTION

- Syntax | 150
- Hierarchy Level | 150
- Description | 150
- Required Privilege Level | 151
- Release Information | 151

Syntax

```
description text-description;
```

Hierarchy Level

```
[edit forwarding-options helpers bootp],
[edit forwarding-options helpers bootp interface (interface-name | interface-group)],
[edit forwarding-options helpers domain],
[edit forwarding-options helpers domain interface interface-name],
[edit forwarding-options helpers port port-number],
[edit forwarding-options helpers port port-number interface interface-name],
[edit forwarding-options helpers tftp],
[edit forwarding-options helpers tftp interface interface-name]
```

Description

Describe a BOOTP, DHCP, Domain Name System (DNS), Trivial File Transfer Protocol (TFTP), or port-based LAN broadcast packet forwarding service, or an interface that is configured for the service.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring DNS and TFTP Packet Forwarding | 112
- Configuring Port-based LAN Broadcast Packet Forwarding | 117
- Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents | 109

dhcp-relay (DHCP Spoofing Prevention)

IN THIS SECTION

- Syntax | 151
- Hierarchy Level | 152
- Description | 152
- Required Privilege Level | 152
- Release Information | 152

Syntax

```
dhcp-relay {
    group group-name {
```

```
        interface interface-name;
    }
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name forwarding-
options],
[edit routing-instances routing-instance-name forwarding-options]
```

Description

Configure Dynamic Host Configuration Protocol (DHCP) snooping on the router. When acting as a snooping agent, the MX Series router typically is located between the client and the DHCP relay agent. It creates filters by “snooping” DHCP messages and binding DHCP-issued IP addresses to the MAC address of the client. These filters help prevent DHCP spoofing.

Configure DHCP snooping by including the appropriate interfaces in the DHCP relay configuration.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4 (MX Series routers only).

RELATED DOCUMENTATION

Preventing DHCP Spoofing on MX Series 5G Universal Routing Platforms | 119

disable (Forwarding Options)

IN THIS SECTION

- Syntax | 153
- Hierarchy Level | 153
- Description | 154
- Required Privilege Level | 154
- Release Information | 154

Syntax

```
disable;
```

Hierarchy Level

```
[edit forwarding-options port-mirror],
[edit forwarding-options port-mirror instance instance-name],
[edit forwarding-options sampling],
[edit forwarding-options sampling instance instance-name],
[edit forwarding-options sampling family  (inet | inet6 | mpls | vpls) ],
[edit forwarding-options sampling family (inet | inet6 | mpls | vpls) output file]
```

Description

Disable traffic accounting, port mirroring, or sampling.

NOTE: The `disable` statement at the `[edit forwarding-options sampling]` hierarchy level disables only Routing Engine-based sampling. To disable PIC-based sampling and inline sampling, include the `disable` statement at the `[edit forwarding-options sampling instance` *`instance-name`*`]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement added to `port-mirror` hierarchy in Junos OS Release 9.6.

NOTE: Beginning in Junos OS Release 15.1F5 and later 15.1 releases and Junos OS Release 16.1 and later, the `disable` option has been deprecated at the `forwarding-options sampling instance` *`instance-name`* `family (inet | inet6 | mpls)` hierarchy level on PTX3000 Series routers. When configured, the option does not take effect, so packets continue to be sampled. Instead of the `disable` option, use the `deactivate forwarding-options sampling instance` *`instance-name`* `family (inet | inet6 | mpls)` command to prevent sampling.

RELATED DOCUMENTATION

Disabling Traffic Sampling | 9

Configuring Traffic Sampling on MX, M and T Series Routers

Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers

domain

IN THIS SECTION

- Syntax | 155
- Hierarchy Level | 155
- Description | 156
- Required Privilege Level | 156
- Release Information | 156

Syntax

```
domain {
    description text-description;
    interface interface-name {
        broadcast;
        description text-description;
        no-listen;
        server address <logical-system logical-system-name> <routing-instance routing-instance-
name>;
    }
    server address <logical-system logical-system-name> <routing-instance routing-instance-name>;
}
```

Hierarchy Level

```
[edit forwarding-options helpers]
```

Description

Enable DNS request packet forwarding.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring DNS and TFTP Packet Forwarding | 112

ecmp-local-bias

IN THIS SECTION

- Syntax | 157
- Hierarchy Level | 157
- Description | 157
- Default | 157
- Required Privilege Level | 157
- Release Information | 158

Syntax

```
ecmp-local-bias;
```

Hierarchy Level

```
[edit forwarding-options load-balance]
```

Description

By default, equal cost multi-path (ECMP) traffic flows are distributed more-or-less equally between forwarding next-hops. Enable `ecmp-local-bias` to have ECMP traffic prefer local forwarding next-hops (that is, local to the packet forwarding engine (PFE) that is performing the packet look up) over remote ones, and to distribute the flows among local members. Local bias percentages cannot be assigned. Note that `ecmp-local-bias` is not intended to be used in conjunction with any other load balancing schemes.

This feature applies chassis-wide. It supports BGP prefixes that are directly reachable with IPv4 MPLS ECMP next-hops on Gigabit Ethernet (ge-) and 10-Gigabit Ethernet (xe-) interfaces. Multicast traffic is not affected.

Use `ecmp-local-bias` when you need to direct ECMP traffic towards local links, for example, to ensure that the overall load on the fabric is reduced.

Default

Not enabled

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1.

RELATED DOCUMENTATION

Configuring Per-Prefix Load Balancing | 85

enhanced-hash-key

IN THIS SECTION

- Syntax | 158
- Hierarchy Level | 160
- Description | 160
- Default | 161
- Options | 161
- Required Privilege Level | 163
- Release Information | 163

Syntax

```
enhanced-hash-key {
    family any {
        incoming-interface-index;
        no-tunnel-payload;
    }
    family inet {
        gtp-tunnel-endpoint-identifier;
        l2tp-tunnel-session-identifier;
        incoming-interface-index;
```

```
        no-destination-port;
        no-source-port;
        type-of-service;
    }
    family inet6 {
        gtp-tunnel-endpoint-identifier;
        incoming-interface-index;
        no-destination-port;
        no-source-port;
        traffic-class;
    }
    family mpls {
        ether-pseudowire {
            zero-control-word;
        }
        incoming-interface-index;
        label-1-exp;
        no-ether-pseudowire;
        no-labels;
        no-payload;
    }
    family multiservice {
        incoming-interface-index;
        no-mac-addresses;
        no-payload;
        outer-priority;
        source-mac;
        no-destination-mac;
    }
    services-loadbalancing {
        family inet {
            layer-3-services {
                destination-address;
                incoming-interface-index;
                source-address;
            }
        }
        family inet6 {
            layer-3-services {
                destination-address;
                incoming-interface-index;
                source-address;
                src-prefix-len;
```

```
            }
        }
    }
    symmetric;
}
```

Hierarchy Level

```
[edit forwarding-options],
[edit logical-systems logical-system-name routing-instances instance-name forwarding-options],
[edit routing-instances instance-name forwarding-options]
```

Description

Computed hash is used in selecting an ECMP path and load balancing. Starting in Junos OS Release 18.3R1, the `flow-label field` is included by default in the hash computation for IPv6, GRE, and PPPoE packets. This can be beneficial, for example, when you have MX routers operating as designated router (DR) or rendezvous point (RP) and want to load balance traffic on the basis of a single Layer 3 or Layer 4 flow. You can revert to the previous method of hash computation by setting the `no-flow-label` option.

- For GRE packets, if the outer IP packet is non-option packet and the inner packet is IPv4 or IPv6, then the source and destination IP addresses from inner packet will be included in hash computation.

  The Layer 4 ports will also be included in hash computation if the protocol of the inner IP packet is TCP or UDP, and if the inner IP packet is not an options packet.

  If outer IP packet is a non-options packet, and the inner packet is MPLS, then the top inner label is included in hash computation.

- For PPPoE packet, if the inner packet is IPv4 or IPv6, then source and destination IP addresses from inner packet will be included in hash computation.

  The Layer 4 ports are included in the hash computation if the protocol of inner IP packet is TCP or UDP, and the inner IP packet is a non-options packet.

For MX Series routers with MPCs, T4000 routers with Type 5 FPCs, EX9200 switches, and PTX10008 routers, select data used in the hash key for enhanced IP forwarding engines.

By default, MPCs use the following parameters for hashing:

- Source IP address
- Destination IP address
- Layer 3 protocol
- Source port
- Destination port
- Generic routing encapsulation (GRE) for GRE packets only.

You can modify the default hashing mechanism on MPCs and Type 5 FPCs by configuring statements at the `[edit forwarding-options enhanced-hash-key]` hierarchy level.

Default

Not enabled.

In PTX, compared to MX (which is similar to QFX), the source and destination mac address options for hash computation are different. While QFX excludes the default MAC address fields for hash computations, PTX includes the destination MAC while excluding the source-mac.

Options

`services-loadbalancing`—Distributes traffic across PICs based on source IP address when a route pointing to more than one services PICs is installed.

`symmetric`—Enable symmetric load balancing across aggregated Ethernet interfaces. This option is needed on Trio-based MPCs only.

Data selections for `services-loadbalancing`:

- `inet`—IPv4 addressing protocol.
- `inet6`—IPv6 addressing protocol.
- `layer-3-services`—Include layer 3 IP data in the hash key.
- `incoming-interface-index`—Include incoming interface index in the hash key.

- source-address—Include source-address in the hash key.
- destination-address—Include destination-address in the hash key.
- src-prefix-len—Include the source prefix length in the hash key.

Data selections for family any:

- incoming-interface-index—(PTX10008 only) Include incoming interface index in the hash key.
- no-tunnel-payload—(PTX10001-36MR, PTX10004, PTX10008, and PTX10016 only) Omit the tunnel payload data from the hash key.

Data selections for family inet:

- **gtp-tunnel-endpoint-identifer**—Include the tunnel endpoint identifier (TEID) field in the hash key for GPRS tunneling protocol (GTP) traffic.

  NOTE: This option is supported only on MX Series routers with MPCs and on the MX80 router.

- incoming-interface-index—Include incoming interface index in the hash key.
- no-destination-port—Omit IP destination port in the hash key.
- no-source-port—Omit IP source port in the hash key.
- type-of-service—Include type-of-service (TOS) byte in the hash key.

Data selections for family inet6:

- **gtp-tunnel-endpoint-identifer**—Include the tunnel endpoint identifier (TEID) field in the hash key for GPRS tunneling protocol (GTP) traffic.

  NOTE: This option is supported only on MX Series routers with MPCs and on the MX80 router.

- incoming-interface-index—Include the incoming interface index in the hash key.
- no-destination-port—Omit the IP destination port in the hash key.
- no-source-port—Omit the IP source port in the hash key.
- traffic-class—Include the traffic class byte in the hash key.

Data selections for family `mpls`:

- `ether-pseudowire`—Load-balance IP over Ethernet pseudowire. Presence of zero-control-word in the payload indicates an Ethernet frame.
- `incoming-interface-index`—Include the incoming interface index in the hash key.
- `label-1-exp`—The EXP bit of the first label is used in the hash calculation.
- `no-ether-pseudowire`—Omit the Ethernet pseudowire payload data from the hash key (MX Series routers with MPCs only).
- `no-labels`—Omit MPLS labels from the hash key (PTX10008 only).
- `no-payload`—Omit the MPLS payload data from the hash key.

Data selections for family `multiservice`:

- `incoming-interface-index`—Include the incoming interface index in the hash key.
- `no-mac-addresses`—Omit source and destination MAC addresses from the hash key.
- `no-payload`—Omit the payload data from the hash key.
- `outer-priority`—Include the outer 802.1 priority bits in the hash key.
- `source-mac`—Includes source MAC address in hash key
- `no-destination-mac`—Excludes destination MAC address in hash key.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.1.

`services-loadbalancing` statement introduced in Junos OS Release 11.2.

`gtp-tunnel-endpoint-identifier` statement introduced in Junos OS Release 13.2

ether-pseudowire statement introduced in Junos OS Release 16.1 for the M Series, MX Series, and PTX Series.

l2tp-tunnel-session-identifier statement introduced in Junos OS Release 17.2

Starting in Junos OS Release 18.3R1, the default behavior for IPv6, GRE, and PPPoE packet hash computation is to include the flow-label field for improved load-balancing in certain cases. Use the no-payload option to revert to the previous method for hash computation.

RELATED DOCUMENTATION

Understanding the Algorithm Used to Load Balance Traffic on MX Series Routers | 61

Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series Routers

export-format

IN THIS SECTION

- Syntax | 164
- Hierarchy Level | 165
- Description | 165
- Options | 165
- Required Privilege Level | 165
- Release Information | 165

Syntax

```
export-format cflowd-version-5;
```

Hierarchy Level

```
[edit forwarding-options monitoring group-name family inet output]
```

Description

Flow monitoring export format.

Options

cflowd-version-5—cflowd version 5.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Passive Flow Monitoring | 35

family (Filtering)

IN THIS SECTION

- Syntax | 166
- Hierarchy Level | 166
- Description | 167
- Options | 167
- Required Privilege Level | 167
- Release Information | 167

Syntax

```
family family-name {
    filter {
         input input-filter-name;
        output output-filter-name;
    }
    flood {
        input filter-name;
    }
    route-accounting;
    source-checking;
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Specify address family for filters.

Options

family-name—Address family. Specify **inet** for IP version 4 (IPv4), **inet6** for IP version 6 (IPv6), **mpls** for MPLS, or **vpls** for virtual private LAN service (VPLS).

> **NOTE**: In Junos OS Release 8.4 and later, the `output` statement is not valid at the `[edit forwarding-options family vpls filter]` hierarchy level.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

route-accounting option introduced in Junos OS Release 8.3; supported only with IPv6.

`source-checking` option introduced in Junos OS Release 12.3 on MX Series 5G Universal Routing Platforms; supported only with IPv6.

RELATED DOCUMENTATION

Applying Forwarding Table Filters | 54

family (Monitoring)

IN THIS SECTION

- Syntax | 168
- Hierarchy Level | 169
- Description | 169
- Required Privilege Level | 169
- Release Information | 169

Syntax

```
family inet {
    output {
        cflowd\ hostname {
            port port-number;
        }
        export-format cflowd-version-5;
        flow-active-timeout seconds;
        flow-export-destination {
            (cflowd-collector | collector-pic);
        }
        flow-inactive-timeout seconds;
        interface interface-name {
            engine-id number;
            engine-type number;
            input-interface-index number;
            output-interface-index number;
            source-address address;
        }
    }
}
```

Hierarchy Level

```
[edit forwarding-options monitoring group-name]
```

Description

Configure flow monitoring for an address family. Only the IPv4 protocol is supported.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Passive Flow Monitoring | 35

family (Port Mirroring)

IN THIS SECTION

Syntax | 170

- Hierarchy Level | 170
- Description | 171
- Options | 171
- Required Privilege Level | 171
- Release Information | 171

Syntax

```
family (any | ccc | inet | inet6 | vpls) {
    output {
        interface interface-name {
            next-hop address;
        }
        next-hop-group group-name{
            group-type inet6;
            interface interface-name {
                next-hop ipv6-address;
            }
            next-hop-subgroup group-name{
                interface interface-name {
                    next-hop ipv6-address;
                }
            }
        }
        no-filter-check;
        server-profile server-profile-name;
    }
}
```

Hierarchy Level

```
[edit forwarding-options port-mirroring],
[edit forwarding-options port-mirroring instance instance-name]
```

Description

Configure the address type family to sample for port mirroring.

Options

`ccc`—Sample Layer 2 VPN traffic.

`inet`—Sample IPv4 traffic.

`inet6`—Sample IPv6 traffic.

`vpls`—Sample VPLS traffic.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`vpls` and `ccc` options introduced in Junos OS Release 9.3 for MX Series routers only.

`vpls` support extended to M7i, M10i, M120, and M320 routers in Junos OS Release 9.5.

`ccc` option introduced in Junos OS Release 9.6 for M120 and M320 routers only.

`ccc` option introduced in Junos OS Release 12.3R2 for EX Series switches.

`next-hop-group` option for `family inet6` introduced in Junos OS Release 14.2 for MX Series routers only.

RELATED DOCUMENTATION

Configuring Port Mirroring | 37

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

family (Sampling)

IN THIS SECTION

- Syntax | 172
- Hierarchy Level | 173
- Description | 173
- Options | 174
- Required Privilege Level | 174
- Release Information | 174

Syntax

```
family (inet | inet6 | mpls) {
    disable;
    output {
        aggregate-export-interval seconds;
        extension-service service-name;
        file {
            disable;
            filename filename;
            files number;
            size bytes;
            (stamp | no-stamp);
            (world-readable | no-world-readable);
        }
        flow-active-timeout seconds;
        flow-inactive-timeout seconds;
        flow-server hostname {
```

```
            aggregation {
                autonomous-system;
                destination-prefix;
                protocol-port;
                source-destination-prefix {
                    caida-compliant;
                }
                source-prefix;
            }
            autonomous-system-type (origin | peer);
            (local-dump | no-local-dump);
            port port-number;
            source-address address;
            version format;
            version9 {
                template template-name;
            }
        }
        interface interface-name {
            engine-id number;
            engine-type number;
            source-address address;
        }
    }
}
```

Hierarchy Level

```
[edit forwarding-options sampling]
```

Description

Configure the protocol family to be sampled.

Options

inet—IP version 4 (IPv4)

inet6—IP version 6 (IPv6)

mpls—MPLS

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

mpls option introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Configuring Traffic Sampling | 6

family inet

IN THIS SECTION

- Syntax | 175
- Hierarchy Level | 175
- Description | 175

- Options | 175
- Required Privilege Level | 176
- Release Information | 176

Syntax

```
family inet {
    layer-3;
    layer-4;
    session-id;
    symmetric-hash {
        complement;
    }
}
```

Hierarchy Level

```
[edit forwarding-options hash-key]
```

Description

Configure layer information for the load-balancing specification. Only the IPv4 protocol is supported.

Options

`family inet`—Incorporate port data into the hash key for flow determination. By default, port data is ignored when determining flows.

- layer-3—Incorporate Layer 3 (IP) data into the hash key. You must include the layer-3 statement. If you omit the layer-3 statement, the management process removes the hash-key statement from the configuration and the router behaves as if you specified layer-3.

  By default, or if you specify only the layer-3 statement, the router uses the following Layer 3 information in the packet header for per-flow load balancing:

  - Source IP address
  - Destination IP address
  - Protocol
  - Incoming interface index

- layer-4—Incorporate Layer 4 Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) data into the hash key.

  If you include the layer-4 statement, the router uses the following Layer 4 information to load-balance:

  - Source port number
  - Destination port number
  - IP type of service

- session-id—Include the session ID in the hash key.
- symmetric-hash—Create the symmetric hash key with source and destination ports.
  - complement—Create the complement of the symmetric hash key.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Per-Packet Load Balancing | 77

family mpls

IN THIS SECTION

- Syntax | 177
- Hierarchy Level | 178
- Description | 178
- Options | 178
- Required Privilege Level | 180
- Release Information | 180

Syntax

```
family mpls {
    all-labels;
    label-1;
    label-2;
    label-3;
    no-labels;
    no-label-1-exp;
    payload {
        ether-pseudowire {
            zero-control-word;
        }
        ip {
            disable;
            layer-3-only;
            port-data {
                source-msb;
                source-lsb;
```

```
                destination-msb;
                destination-lsb;
            }
        }
    }
}
```

Hierarchy Level

```
[edit forwarding-options hash-key]
```

Description

For aggregated Ethernet and SONET/SDH interfaces only, configure load balancing based on MPLS labels and payload. Only the IPv4 protocol is supported.

Options

`family mpls`—(Aggregated Ethernet interfaces, aggregated SONET/SDH interfaces, and multiple equal-cost MPLS next hops only) Incorporate MPLS label and payload information into the hash key for per-flow load balancing. Only the IPv4 protocol is supported.

- `all-labels`—(PTX Series Packet Transport Routers only) Up to eight MPLS labels are included in the hash key to identify the uniqueness of a flow in the Packet Forwarding Engine. This is the default setting.
- `label-1`—(M120, M320, MX Series, and T Series routers only) Include the first MPLS label into the hash key. This is used for a one-label packet for per-flow load balancing IPv4 VPLS traffic based on IP information and MPLS labels.
- `label-2`—(M120, M320, MX Series, and T Series routers only) Include the second MPLS label into the hash key. This is used for a two-label packet for per-flow load balancing IPv4 VPLS traffic based on IP information and MPLS labels. To use the second MPLS label in the hash key, include both the `label-1` and `label-2` statements at the `[edit forwarding-options hash-key family mpls]` hierarchy level. By default,

the router provides hashing on the first and second labels. If both labels are specified, the entire first label and the first 16 bits of the second label are hashed.

- label-3—(M120, M320, MX Series, and T Series routers only) Include the third MPLS label into the hash key. To use the third MPLS label, include the label-1, label-2, and label-3 statements at the [edit forwarding-options hash-key family mpls] hierarchy level.
- no-labels—Include no MPLS labels into the hash key.
- no-label-1-exp—(M120, M320, MX Series, and T Series routers only) The EXP bit of the first label is not used in the hash calculation to avoid reordering complications.
- payload—Incorporate bits from the IP payload into the hash key for per-flow load balancing Layer 2 information based on MPLS labels.
  - disable—(PTX Series Packet Transport Routers only) Exclude IP payload from the hash key.
  - ether-pseudowire—(M120, M320, MX Series, and T Series routers only) Load-balance IPv4 traffic over Layer 2 Ethernet pseudowires.
  - zero-control-word—(M Series, MX Series, and PTX Series) Precedes Ethernet packet to indicate the start of an Ethernet frame in an MPLS ether-pseudowire payload.
  - ip—Include the IP address of the IPv4 or IPv6 payload into the hash key for per-flow load balancing Layer 2 information based on MPLS labels. For the PTX Series Packet Transport Routers, this is the default setting with both Layer 3 and Layer 4 IP information included in the hash key.
    - disable—(PTX Series Packet Transport Routers only) Exclude IP payload from the hash key.
    - layer-3-only—Include only Layer 3 IP information from the IP payload data into the hash key for per-flow load balancing Layer 2 information based on MPLS labels.
    - port-data—(M120, M320, MX Series, and T Series routers only) Include the source and destination port field information into the hash key. By default, the most significant byte and least significant byte of the source and destination port fields are hashed. To select specific bytes to be hashed, include one or more of the source-msb, source-lsb, destination-msb, and destination-lsb options at the [edit forwarding-options hash-key family mpls payload ip port-data] hierarchy level. To prevent all four bytes from being hashed, include the layer-3-only statement at the [edit forwarding-options hash-key family mpls payload ip] hierarchy level.
      - destination-lsb—Include the least-significant byte of the destination port.
      - destination-msb—Include the most-significant byte of the destination port.
      - source-lsb—Include the least-significant byte of the source port.

- source-msb—Include the most-significant byte of the source port.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

no-label-1-exp option introduced in Junos OS Release 8.0.

label-3 and no-labels options introduced in Junos OS Release 8.1.

ether-pseudowire statement introduced in Junos OS Release 9.1 (M320 and T Series routers only); support extended to M120 and MX Series routers in Junos OS Release 9.4.

all-labels and payload ip disable options introduced in Junos OS Release 12.1X48R2. (PTX Series Packet Transport Routers only).

zero-control-word option introduced in Junos OS Release 16.1 for the M Series, MX Series, and PTX Series.

RELATED DOCUMENTATION

Configuring Load Balancing Based on MPLS Labels

Configuring Load Balancing for Ethernet Pseudowires

family multiservice

IN THIS SECTION

- Syntax | 181
- Hierarchy Level | 182
- Description | 182
- Options | 182
- Required Privilege Level | 183
- Release Information | 183

Syntax

```
family multiservice {
    destination-mac;
    label-1;
    label-2;
    payload {
        ip {
            layer-3 {
                (source-ip-only | destination-ip-only);
            }
            layer-3-only;
            layer-4;
        }
    }
    source-mac;
    symmetric-hash {
        complement;
    }
}
```

Hierarchy Level

```
[edit forwarding-options hash-key]
```

Description

Configure load balancing based on Layer 2 media access control information. On MX Series routers, configure VPLS load balancing. On M120 and M320 routers only, configure VPLS load balancing based on MPLS labels and IP information. For IPv4 traffic, only the IP source and destination addresses are included in the hash key. For MPLS and IPv4 traffic, one or two MPLS labels and IPv4 source and destination addresses are included. For MPLS Ethernet pseudowires, only one or two MPLS labels are included in the hash key.

Options

You can configure one or more options to load-balance using the packet information that you specify.

`destination-mac`—Include the destination-address MAC information in the hash key for Layer 2 load balancing.

`label-1` (M120 and M320 routers only)—Include the first MPLS label in the hash key. Used for including a one-label packet for per-flow load balancing of IPv4 VPLS traffic based on IP information and MPLS labels.

`label-2` (M120 and M320 routers only)—Include the second MPLS label in the hash key. If both `label-1` and `label-2` are specified, the entire first label and the first 16 bits of the second label are hashed.

`payload` (MX Series, M120, and M320 routers only)—Include the packet's IP payload in the hash key.

- `ip` (MX Series, M120, and M320 routers only)—Include the IP address of the IPv4 or IPv6 payload in the hash key.
  - `layer-3` (MX Series routers only)—Use this to include Layer 3 information from the packet's IP payload in the hash key.
    - `destination-ip-only` (MX Series routers only)—Use this to include only the destination IP address in the payload in the hash key.

- source-ip-only (MX Series routers only)—Use this to include only the source IP address in the payload in the hash key.

  NOTE: You can include either the source-ip-only or the destination-ip-only statement, not both. They are mutually exclusive.

- layer-3-only (M120, and M320 routers only)—Include only the Layer 3 information from the packet's IP payload in the hash key.
- layer-4 (MX Series routers only)—Include Layer 4 information from the packet's IP payload in the hash key.

  NOTE: On MX Series routers only, you can configure either Layer 3 or Layer 4 load balancing, or both at the same time.

  NOTE: On I chip platforms, an unknown Layer 4 header is excluded from load-balance hashing to avoid undesired packet reordering.

source-mac—Include the source-address MAC information in the hash key.

symmetric-hash (MX Series routers only)—Configure the symmetric hash or symmetric hash complement for configuring symmetrical load balancing on an 802.3ad Link Aggregation Group.

- complement —Include the complement of the symmetric hash in the hash key.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

ip, label-1, label-2, layer-3-only, and payload options introduced in Junos OS Release 9.4.

layer-3, layer-. source-ip-only, and destination-ip-only options introduced in Junos OS Release 9.5.

symmetric-hash and complement options introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

- Configuring Load Balancing Based on MAC Addresses
- *Configuring VPLS Load Balancing Based on IP and MPLS Information*
- *Configuring VPLS Load Balancing on MX Series 5G Universal Routing Platforms*
- *Configuring VPLS Load Balancing*

file (Extended DHCP Relay Agent and Helpers Trace Options)

IN THIS SECTION

- Syntax | 184
- Hierarchy Level | 185
- Description | 185
- Options | 185
- Required Privilege Level | 185
- Release Information | 185

Syntax

```
file filename  <files number> <match regular-expression> <size bytes> <world-readable | no-world-readable>;
```

Hierarchy Level

```
[edit forwarding-options dhcp-relay traceoptions],
[edit forwarding-options helpers traceoptions]
```

Description

Configure information about the DNS and TFTP packet-forwarding files that contain trace logging information.

Options

filename—Name of the file containing the trace information.

- **Default: /var/log/sampled**

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring DNS and TFTP Packet Forwarding | 112

file (Sampling)

IN THIS SECTION

- Syntax | 186
- Hierarchy Level | 186
- Description | 186
- Required Privilege Level | 187
- Release Information | 187

Syntax

```
file filename filename  <disable> <files number> <stamp | no-stamp> <size bytes> <world-readable | no-world-readable>;
```

Hierarchy Level

```
[edit forwarding-options sampling family family-name output]
```

Description

Collect the traffic samples in a file.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Collecting Traffic Sampling Output in a File | 10

file (Trace Options)

IN THIS SECTION

- Syntax | 187
- Hierarchy Level | 188
- Description | 188
- Options | 188
- Required Privilege Level | 188
- Release Information | 188

Syntax

```
file filename  <files number> <size bytes> <world-readable | no-world-readable>;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring traceoptions],
[edit forwarding-options sampling traceoptions]
```

Description

Configure information about the files that contain trace logging information.

Options

filename—The name of the file containing the trace information.

- **Default: /var/log/sampled**

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Tracing Traffic-Sampling Operations | 22

filename (Sampling)

IN THIS SECTION

- Syntax | 189
- Hierarchy Level | 189
- Description | 189
- Options | 190
- Required Privilege Level | 190
- Release Information | 190

Syntax

```
filename filename;
```

Hierarchy Level

```
[edit forwarding-options sampling family family-name output file]
```

Description

Configure the name of the output file.

Options

filename—Name of the file in which to place the traffic samples. All files are placed in the directory **/var/tmp**.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Collecting Traffic Sampling Output in a File | 10

files (Sampling and Traceoptions)

IN THIS SECTION

- Syntax | 191
- Hierarchy Level | 191
- Description | 191
- Options | 191
- Required Privilege Level | 191
- Release Information | 192

Syntax

```
files number;
```

Hierarchy Level

```
[edit forwarding-options helperstraceoptions file],
[edit forwarding-options port-mirroring traceoptions file],
[edit forwarding-options sampling family family-name output file],
[edit forwarding-options sampling traceoptions file]
```

Description

Configure the total number of files to be saved with samples or trace data.

Options

number—Maximum number of traffic sampling or trace log files. When a file named ***sampling-file*** reaches its maximum size, it is renamed ***sampling-file*.0**, then ***sampling-file*.1**, and so on, until the maximum number of traffic sampling files is reached. Then the oldest sampling file is overwritten.

- **Range:** 1 through 100 files
- **Default:** 5 files for sampling output; 10 files for trace log information

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Collecting Traffic Sampling Output in a File | 10

Tracing Traffic-Sampling Operations | 22

filter (IPv4, IPv6, and MPLS)

IN THIS SECTION

- Syntax | 192
- Hierarchy Level | 193
- Description | 193
- Options | 193
- Required Privilege Level | 193
- Release Information | 193

Syntax

```
filter {
    input input-filter-name;
    output output-filter-name;
}
```

Hierarchy Level

```
[edit forwarding-options family (inet | inet6 | mpls)]
```

Description

Apply a forwarding table filter to a forwarding table.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying Forwarding Table Filters | 54

filter (VPLS)

IN THIS SECTION

- Syntax | 194
- Hierarchy Level | 194
- Description | 194
- Options | 195
- Required Privilege Level | 195
- Release Information | 195

Syntax

```
filter input filter-name;
```

Hierarchy Level

```
[edit forwarding-options family vpls],
[edit logical-systems logical-system-name forwarding-optionsfamily vpls],
[edit logical-systems logical-system-name routing -instances instance-name forwarding-
optionsfamily vpls],
[edit routing -instances instance-name forwarding-optionsfamily vpls]
```

Description

Apply a forwarding table filter for VPLS.

Options

The other statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying Forwarding Table Filters | 54

flood

IN THIS SECTION

- Syntax | 196
- Hierarchy Level | 196
- Description | 196
- Options | 196
- Required Privilege Level | 196
- Release Information | 196

Syntax

```
flood {
    input filter-name;
}
```

Hierarchy Level

```
[edit forwarding-options family vpls],
[edit routing -instances instance-name forwarding-options family vpls]
```

Description

Apply a forwarding table filter to a flood table.

Options

input *filter-name*—Name of the forwarding table filter.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying Forwarding Table Filters | 54

Layer 2 Port Mirroring Firewall Filters

flow-active-timeout

IN THIS SECTION

- Syntax | 197
- Hierarchy Level | 197
- Description | 198
- Options | 198
- Required Privilege Level | 198
- Release Information | 198

Syntax

```
flow-active-timeout seconds;
```

Hierarchy Level

```
[edit forwarding-options accounting group-name output],
[edit forwarding-options monitoring group-name family inet output ],
[edit forwarding-optionssampling family family-name output]
```

Description

Configure the time that elapses before another active flow is exported.

Options

seconds—Timeout, in seconds.

- **Range:** 60 through 1800
- **Default:** 1800

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Discard Accounting | 31

Configuring Passive Flow Monitoring | 35

Collecting Traffic Sampling Output in a File | 10

flow-export-destination

IN THIS SECTION

- Syntax | 199
- Hierarchy Level | 199
- Description | 199
- Options | 200
- Required Privilege Level | 200
- Release Information | 200

Syntax

```
flow-export-destination {
    (cflowd-collector | collector-pic);
}
```

Hierarchy Level

```
[edit forwarding-options monitoring group-name family inet output]
```

Description

Configure flow collection.

Options

cflowd-collector—cflowd collector.

collector-pic—Collector PIC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Passive Flow Monitoring | 35

flow-inactive-timeout

IN THIS SECTION

- Syntax | 201
- Hierarchy Level | 201
- Description | 201
- Options | 201
- Required Privilege Level | 201
- Release Information | 202

Syntax

```
flow-inactive-timeout seconds;
```

Hierarchy Level

```
[edit forwarding-options accounting group-name output],
[edit forwarding-options monitoring group-name family inet output],
[edit forwarding-options sampling family family-name output]
```

Description

Configure the time that elapses before a flow is considered inactive.

Options

seconds—Timeout, in seconds.

- **Range:** 15 through 1800
- **Default:** 60

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Discard Accounting | 31

Configuring Passive Flow Monitoring | 35

Collecting Traffic Sampling Output in a File | 10

flow-server

IN THIS SECTION

- Syntax | 202
- Hierarchy Level | 203
- Description | 203
- Options | 203
- Required Privilege Level | 204
- Release Information | 204

Syntax

```
flow-server hostname {
    aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
            caida-compliant;
        }
```

```
        source-prefix;
    }
    autonomous-system-type (origin | peer);
    (local-dump | no-local-dump);
    port port-number;
    routing-instance instance-name;
    source-address address;
    version9 {
        template template-name;
    }
    version-ipfix {
        template template-name;
    }
}
```

Hierarchy Level

```
[edit forwarding-options sampling family family-name output]
```

Description

Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility **cfdcollect**. Specify a host system to collect sampled flows using the version 9 format.

You can configure up to one version 5 and one version 8 flow format at the `[edit forwarding-options sampling output flow-server hostname]` hierarchy level. For the same configuration, you can specify only either version 9 flow record formats or formats using versions 5 and 8, not both types of formats.

Options

hostname—The IP address or identifier of the host system (the workstation either running the cflowd utility or collecting traffic flows using version 9).

You can configure only one host system for version 9.

On QFX10002 switches, only IPv4 host systems are supported.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`version9` statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Collecting Traffic Sampling Output in the Cisco Systems NetFlow Services Export Version 9 Format | 16

Directing Traffic Sampling Output to a Server Running the cflowd Application | 12

group (DHCP Spoofing Prevention)

IN THIS SECTION

- Syntax | 205
- Hierarchy Level | 205
- Description | 205
- Required Privilege Level | 205
- Release Information | 205

Syntax

```
group group-name {
    interface interface-name;
}
```

Hierarchy Level

```
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name forwarding-options
dhcp-relay],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay]
```

Description

Configure Dynamic Host Configuration Protocol (DHCP) snooping on the router. When acting as a snooping agent, the MX Series router typically is located between the client and the DHCP relay agent. It creates filters by “snooping” DHCP messages and binding DHCP-issued IP addresses with the MAC address of the client. These filters help prevent DHCP spoofing.

Configure DHCP snooping by including the appropriate interfaces under the `group` statement.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4 (MX Series routers only).

RELATED DOCUMENTATION

Preventing DHCP Spoofing on MX Series 5G Universal Routing Platforms | 119

gtp-tunnel-endpoint-identifier

IN THIS SECTION

- Syntax | 206
- Hierarchy Level | 206
- Hierarchy Level (QFX5000 line of switches) | 207
- Hierarchy Level (QFX10000 line of switches) | 207
- Description | 207
- Required Privilege Level | 208
- Release Information | 208

Syntax

```
gtp-tunnel-endpoint-identifier;
```

Hierarchy Level

```
[edit forwarding-options hash-key family inet layer-4],
```

```
[edit forwarding-options hash-key family inet6 layer-4]
```

Hierarchy Level (QFX5000 line of switches)

```
[edit forwarding-options enhanced-hash-key family inet]
```

Hierarchy Level (QFX10000 line of switches)

```
[edit forwarding-options enhanced-hash-key family inet],
```

```
[edit forwarding-options enhanced-hash-key family inet6]
```

Description

When you configure `gtp-tunnel-endpoint-identifier`, the hash calculation of IPv4 or IPv6 packets are included in the GPRS tunneling protocol–tunnel endpoint ID (GTP-TEID) field hash calculations.

NOTE: The `gtp-tunnel-endpoint-identifier` configuration statement is supported on PTX Series routers only when network services is set to `enhanced-mode`. For more information, see `enhanced-mode`.

On the QFX5000 and QFX10000 lines of switches, if the `gtp-tunnel-endpoint-identifier` statement is configured, the default Layer 4 port 2152 (and 2123 as well on QFX5000) is set to use along with the default first byte 0x32.

(On the QFX5000 line of switches only) In most cases, configuring the `gtp-tunnel-endpoint-identifier` statement on QFX5000 switches is sufficient for enabling GTP hashing. After you have enabled GTP hashing, if GTP hashing does not work, we recommend that you capture the packets using relevant tools and identify the offset value. As per standards, 0x32 is the default header offset value. But, due to some special patterns in the header, offset may vary to say 0x30, 0x28, and so on. In these cases, use the `gtp-header-offset` statement to set a proper offset value. Once the header offset value is resolved, run the `gtp-tunnel-endpoint-identifier` command for enabling GTP hashing successfully. Refer to *gtp-header-offset* for more details.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F3 and 16.1R2.

RELATED DOCUMENTATION

hash-key (Forwarding Options) | 208

Understanding Per-Packet Load Balancing | 75

Configuring Per-Packet Load Balancing | 77

gtp-header-offset

hash-key (Forwarding Options)

IN THIS SECTION

- Syntax | 209
- Hierarchy Level | 210
- Description | 210
- Required Privilege Level | 211
- Release Information | 211

Syntax

```
hash-key {
    family inet {
        layer-3;
        layer-4;
        session-id;
        symmetric-hash {
            complement;
        }
    }
    family mpls {
        all-labels;
        bottom-label-1;
        bottom-label-2;
        bottom-label-3;
        label-1;
        label-2;
        label-3;
        no-labels;
        no-label-1-exp;
        payload {
            ether-pseudowire {
                zero-control-word;
            }
            ip {
                disable;
                layer-3-only;
                port-data {
                    destination-lsb;
                    destination-msb;
                    source-lsb;
                    source-msb;
                }
            }
        }
    }
    family multiservice {
        destination-mac;
        label-1;
        label-2;
```

```
        payload {
            ip {
                layer-3-only;
                layer-3 {
                    (source-ip-only | destination-ip-only);
                }
                layer-4;
            }
        }
        source-mac:
    }
}
```

Hierarchy Level

```
[edit forwarding-options
```

Description

Select which packet header data to use for per-flow load balancing.

The options are explained separately.

NOTE: To modify the default hashing mechanism on Modular Port Concentrators (MPCs) and Type 5 FPCs, you need to configure the statements at the `[edit forwarding-options enhanced-hash-key]` hierarchy level. Statements at the `[edit forwarding-options hash-key]` hierarchy level do not support MPCs and Type 5 FPCs.

NOTE: The following statements are not supported on T Series routers:

- The `symmetric-hash` and the `session-id` statements at the `[edit forwarding-options hash-key family inet]` hierarchy level and all statements at the `[edit forwarding-options hash-key family multiservice]` hierarchy level.
- The `label-1` and `label-2` statements, and the IP configuration at the `[edit forwarding-options hash-key family multiservice]` hierarchy level.

NOTE: The following statements are not supported on Q Series switches:

- The `symmetric-hash` and the `session-id` statements at the `[edit forwarding-options hash-key family inet]` hierarchy level.

NOTE: On PTX platforms, *family inet/inet6 layer-3 traffic-class (dscp-v4/dscp-v6)* is disabled by default.

NOTE: EVPN-VXLAN configuration is required for underlay ECMP to work in ACX platforms on Junos EVO.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`family multiservice` and `no-label-1-exp` options introduced in Junos OS Release 8.0.

`label-3` and `no-labels` options introduced in Junos OS Release 8.1.

`ether-pseudowire` statement introduced in Junos OS Release 9.1 (M320 and T Series routers only); support extended to M120 and MX Series routers in Junos OS Release 9.4.

ip, label-1, label-2, layer-3-only, and payload options for the family multiservice statement introduced in Junos OS Release 9.4 (M120 and M320 routers only). For MX Series routers, only the ip and payload statements apply.

layer-3, source-ip-only, destination-ip-only, and layer-4 statements introduced for the family multiservice statement in Junos OS Release 9.5. (MX Series routers only).

all-labels and payload ip disable statements introduced for the family mpls statement in Junos OS Release 12.1X48R2 PTX Series Packet Transport Routers only.

bottom-label statement introduced for the family mpls statement in Junos OS Release 14.1 for MX Series routers with DPCs (excluding M7i, M10i, and M120.)

zero-control-word option for the ether-pseudowire statement introduced in Junos OS Release 16.1 for the M Series, MX Series, and PTX Series.

RELATED DOCUMENTATION

Configuring Per-Packet Load Balancing | 77

Configuring Load Balancing Based on MPLS Labels

Configuring Load Balancing Based on MAC Addresses

helpers

IN THIS SECTION

- Syntax | 213
- Hierarchy Level | 214
- Description | 215
- Required Privilege Level | 216
- Release Information | 216

Syntax

```
helpers {
    bootp {
        client-response-ttl number;
        description text-description;
        interface interface-group {
            client-response-ttl number;
            description text-description;
            maximum-hop-count number;
            minimum-wait-time seconds;
            no-listen;
            server address {
                logical-system logical-system-name <routing-instance [ <default> routing-
instance-names ]>;
                routing-instance [ <default> routing-instance-names ];
            }
        }
        maximum-hop-count number;
        minimum-wait-time seconds;
        relay-agent-option;
        server address {
            logical-system logical-system-name <routing-instance [ <default> routing-instance-
names ]>;
            routing-instance [ <default> routing-instance-names ];
        }
    }
    domain {
        description text-description;
        interface interface-name {
            broadcast;
            description text-description;
            no-listen;
            server address <logical-system logical-system-name> <routing-instance routing-
instance-name>;
        }
        server address <logical-system logical-system-name> <routing-instance routing-instance-
name>;
    }
    port (Packet Forwarding) port-number {
        description text-description;
```

```
        interface interface-name {
            broadcast;
            description text-description;
            no-listen;
            server address <logical-system logical-system-name> <routing-instance routing-
instance-name>;
        }
        server address <logical-system logical-system-name> <routing-instance routing-instance-
name>;
    }
    tftp {
        description text-description;
        interface interface-name {
            broadcast;
            description text-description;
            no-listen;
            server address <logical-system logical-system-name> <routing-instance routing-
instance-name>;
        }
        server address <logical-system logical-system-name> <routing-instance routing-instance-
name>;
    }
    traceoptions {
        file filename  <files number> <match regular-expression> <size bytes> <world-readable | no-
world-readable>;
        flag flag;
        level level;
        no-remote-trace level;
    }
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Enable TFTP or DNS request packet forwarding, or configure the router, switch, or interface to act as a DHCP/BOOTP relay agent. Use only one server address per interface or global configuration.

You can also use the helpers `port` statement to enable forwarding LAN broadcast traffic on custom UDP ports to particular servers as unicast traffic. Configure the UDP port number and optionally an interface on which to listen for broadcast traffic, and the destination server address to receive that traffic, as shown in either of the following sample configurations:

```
user@ host# show forwarding-options
helpers {
    port 1200 {
        server 10.20.30.40;
        }
    }
```

```
user@ host# show forwarding-options
helpers {
    port 3000 {
        interface {
            fe-0/0/1.0 {
                server 192.0.2.2;
            }
        }
    }
    port 3001 {
        interface {
            fe-0/0/0.0 {
                server 192.0.2.2;
            }
        }
    }
}
```

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring DNS and TFTP Packet Forwarding | 112

Configuring Port-based LAN Broadcast Packet Forwarding | 117

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents | 109

hosted-service-identifier

IN THIS SECTION

- Syntax | 217
- Hierarchy Level | 217
- Description | 217
- Options | 217
- Required Privilege Level | 217
- Release Information | 217

Syntax

```
hosted-service-identifier identifier;
```

Hierarchy Level

```
[edit services hosted-services server-profile server-profile-name]
```

Description

Configure the identifier for the service performed on the remote server

Options

identifier—Identifer for the service performed on the remote server.

- **Range:** 1 through 63

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

hosted-services

IN THIS SECTION

- Syntax | 218
- Hierarchy Level | 218
- Description | 219
- Options | 219
- Required Privilege Level | 219
- Release Information | 219

Syntax

```
hosted-services {
    server-profile server-profile-name {
        client-address ipv4-address;
        server-address ipv4-address;
    }
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure services performed on the remote server.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

hyper-mode (forwarding-options)

IN THIS SECTION

- Syntax | 220
- Hierarchy Level | 220
- Description | 220

Required Privilege Level | 221
Release Information | 221

Syntax

```
(hyper-mode | no-hyper-mode);
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Configure the hyper mode feature to increase the rate at which a data packet is processed. This configuration results in the optimization of the lifetime of a data packet, which further enables the network device (a router or a switch) to provide better performance and throughput. This feature is supported on enhanced MPCs such as MPC3E, MPC4E, MPC5E, and MPC6E.

Configure the `no-hyper-mode` statement when you want to turn off hyper mode and have the device operate in normal mode.

NOTE: The hyper mode feature is configured at the global level and requires a system reboot to turn it on or off. You can enable the feature only if the network-service mode on the device is configured as either `enhanced-ip` or `enhanced-ethernet`.

"Understanding the Hyper Mode Feature on Enhanced MPCs for MX Series Routers and EX9200 Switches" on page 121 contains a list of features that are not supported in hyper mode.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3R4.

RELATED DOCUMENTATION

Configuring Hyper Mode on Enhanced MPCs to Speed Up Packet Processing | 124

Understanding the Hyper Mode Feature on Enhanced MPCs for MX Series Routers and EX9200 Switches | 121

show forwarding-options hyper-mode | 355

indexed-load-balance

IN THIS SECTION

- Syntax | 222
- Hierarchy Level | 222
- Description | 222
- Default | 222
- Required Privilege Level | 223
- Release Information | 223

Syntax

```
indexed-load-balance;
```

Hierarchy Level

```
[edit forwarding-options load-balance],
[edit logical-systems logical-system-name forwarding-options load-balance],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options load-balance],
[edit routing-instances routing-instance-name forwarding-options load-balance]
```

Description

The indexed-load-balance statement causes the creation of a nexthop structure that is both a function of the hash, and a function of the low-order bits of the IP address. Include this statement if you notice issues with load-balance distribution for IPv4 or IPv6 traffic to improve load-balance distribution for unicast and aggregated next hops.

For MPC line cards in MX routers, indexed-load-balance has been superseded by an internal hash-rotation mechanism to reduce polarization.

CAUTION: Including the indexed-load-balance statement causes an increase in memory usage on the device.

Default

Disabled

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Starting with Junos OS Release 12.1, the indexed-next-hop statement has been renamed as the indexed-load-balance statement.

RELATED DOCUMENTATION

Understanding the Algorithm Used to Load Balance Traffic on MX Series Routers | 61

Configuring Per-Prefix Load Balancing | 85

enhanced-hash-key | 158

input (Forwarding Table)

IN THIS SECTION

- Syntax | 224
- Hierarchy Level | 224
- Description | 224
- Options | 224
- Required Privilege Level | 224
- Release Information | 224

Syntax

```
input filter-name;
```

Hierarchy Level

```
[edit forwarding-options family (inet | inet6 | mpls | vpls) filter],
[edit routing-instances routing-instance-name forwarding-options family (inet | inet6 | mpls |
vpls) filter]
```

Description

Apply a forwarding table filter to ingress traffic of the forwarding table.

Options

filter-name—Name of the applied filter.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying Forwarding Table Filters | 54

input (Port Mirroring)

IN THIS SECTION

- Syntax | 225
- Hierarchy Level | 225
- Description | 226
- Required Privilege Level | 226
- Release Information | 226

Syntax

```
input {
    maximum-packet-length bytes;
    rate number;
    run-length number;
}
```

Hierarchy Level

```
[edit forwarding-options port-mirroring],
[edit forwarding-options port-mirroring instance instance-name]
```

Description

Configure input packet properties for port mirroring.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

maximum-packet-length option introduced in Junos OS Release 9.6 for M120 and M320 routers only.

RELATED DOCUMENTATION

Configuring Port Mirroring | 37

input (Sampling)

IN THIS SECTION

- Syntax | 227
- Hierarchy Level | 227
- Description | 227
- Required Privilege Level | 227
- Release Information | 227

Syntax

```
input {
    max-packets-per-second number;
    maximum-packet-length bytes;
}
```

Hierarchy Level

```
[edit forwarding-options sampling ]
```

Description

Configure traffic sampling on a logical interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for sampling of MPLS traffic introduced in Junos OS Release 8.3.

instance

IN THIS SECTION

- Syntax | 228
- Hierarchy Level | 229
- Description | 229
- Options | 229
- Required Privilege Level | 229
- Release Information | 229

Syntax

```
instance {
    instance-name {
        input {
            maximum-packet-length bytes;
            rate number;
            run-length number;
        }
        family (any | ccc| inet | inet6 | mpls | vpls) {
            output {
                interface interface-name {
                    next-hop address;
                }
                no-filter-check;
                server-profile server-profile-name;
            }
        }
    }
}
```

Hierarchy Level

```
[edit forwarding-options port-mirroring],
[edit routing-instances routing-instance-name forwarding-options port-mirroring]
```

Description

Configure a port-mirroring instance.

Options

port-mirroring-instance-name—Name of the port-mirroring instance.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3 (MX Series routers only).

Support extended to M120 and M320 routers in Junos OS Release 9.5.

maximum-packet-length and **ccc** options introduced in Junos OS Release 9.6 for M120 and M320 routers only.

`server-profile` option introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers only.

RELATED DOCUMENTATION

Configuring Port Mirroring | 37

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

interface (Accounting or Sampling)

IN THIS SECTION

- Syntax | 230
- Hierarchy Level | 230
- Description | 231
- Options | 231
- Required Privilege Level | 231
- Release Information | 231

Syntax

```
interface interface-name {
    engine-id number;
    engine-type number;
    source-address address;
}
```

Hierarchy Level

```
[edit forwarding-options accounting group-name output],
[edit forwarding-options sampling family family-name output]
```

Description

Specify the output interface for sending copies of packets elsewhere to be analyzed.

Options

engine-id *number*—Identity of the accounting interface.

engine-type *number*—Type of this accounting interface.

interface-name—Name of the accounting interface.

source-address *address*—Address used for generating packets.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Discard Accounting | 31

Collecting Traffic Sampling Output in a File | 10

interface (BOOTP)

IN THIS SECTION

- Syntax | 232
- Hierarchy Level | 232
- Description | 233
- Options | 233
- Required Privilege Level | 233
- Release Information | 233

Syntax

```
interface (interface-name | interface-group) {
    apply-secondary-as-giaddr;
    broadcast;
    client-response-ttl number;
    description text-description;
    maximum-hop-count number;
    minimum-wait-time seconds;
    no-listen;
    server address {
        logical-system logical-system-name <routing-instance [ <default> routing-instance-
names ]>;
        routing-instance [ <default> routing-instance-names ];
    }
}
```

Hierarchy Level

```
[edit forwarding-options helpers bootp]
```

Description

Specify the interface for a DHCP and BOOTP relay agent.

Options

interface-group	Set a logical interface or group of logical interfaces with a specific DHCP relay configuration.
apply-secondary-as-giaddr	Enable DHCP relay to use secondary gateway IP on this interface.
broadcast	If the layer 2 interface is unknown, then broadcast.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents | 109

interface (DHCP Spoofing Prevention)

IN THIS SECTION

- Syntax | 234
- Hierarchy Level | 234
- Description | 234
- Required Privilege Level | 235
- Release Information | 235

Syntax

```
interface interface-name;
```

Hierarchy Level

```
[edit routing-instances routing-instance-name bridge-domains bridge-domain-name forwarding-options
dhcp-relay group group-name],
[edit routing-instances routing-instance-name forwarding-options dhcp-relay group group-name]
```

Description

Configure Dynamic Host Configuration Protocol (DHCP) snooping on the router. When acting as a snooping agent, the MX Series router typically is located between the client and the DHCP relay agent. It creates filters by “snooping” DHCP messages and binding DHCP-issued IP addresses with the MAC address of the client. These filters help prevent DHCP spoofing.

DHCP snooping is configured by including the appropriate interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4 (MX Series routers only).

RELATED DOCUMENTATION

Preventing DHCP Spoofing on MX Series 5G Universal Routing Platforms | 119

interface (DNS, Port, and TFTP Packet Forwarding or Relay Agent)

IN THIS SECTION

- Syntax | 236
- Hierarchy Level | 236
- Description | 236
- Options | 236
- Required Privilege Level | 237
- Release Information | 237

Syntax

```
interface interface-name {
    broadcast;
    description text-description;
    no-listen;
    server address <logical-system logical-system-name> <routing-instance routing-instance-name>;
}
```

Hierarchy Level

```
[edit forwarding-options helpers domain],
[edit forwarding-options helpers port port-number],
[edit forwarding-options helpers tftp]
```

Description

Specify the interface for monitoring and forwarding DNS or TFTP requests, or for forwarding LAN broadcast traffic on a custom UDP port to a particular server as unicast traffic.

When configuring port helpers, in releases prior to Junos OS Release 17.2, only one server can be specified for a given port. For Junos OS Release 17.2 and later, multiple servers can be specified for a given port at the global or interface-specific level. In this case, the same packet, with the originator IP address and port requests, is forwarded to the different configured servers; the payload of the UDP packet is not modified.

Options

interface-name—Name of the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced for port helpers in Junos OS Release 17.2R1 for EX4300 switches.

Support for multiple server instances for a given port introduced in Junos OS Release 17.2 for MX Series routers.

Support for multiple server instances for a given port introduced in Junos OS Release 17.3R1 for EX9200 switches.

RELATED DOCUMENTATION

Configuring DNS and TFTP Packet Forwarding | 112

Configuring Port-based LAN Broadcast Packet Forwarding | 117

interface (Monitoring)

IN THIS SECTION

- Syntax | 238
- Hierarchy Level | 238
- Description | 238
- Options | 238
- Required Privilege Level | 239
- Release Information | 239

Syntax

```
interface interface-name {
    engine-id number;
    engine-type number;
    input-interface-index number;
    output-interface-index number;
    source-address address;
}
```

Hierarchy Level

```
[edit forwarding-options monitoring group-name inet output]
```

Description

Specify the output interface for monitored traffic.

Options

interface-name—Name of the interface.

engine-id *number*—Identity of the monitoring interface.

engine-type *number*—Type of the monitoring interface.

input-interface-index *number*—Input interface index for records from the interface.

output-interface-index *number*—Output interface index for records from the interface.

source-address *address*—Address used for generating packets.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Passive Flow Monitoring | 35

interface (Next-Hop Group)

IN THIS SECTION

- Syntax | 240
- Hierarchy Level | 240
- Description | 240
- Options | 240
- Required Privilege Level | 240
- Release Information | 241

Syntax

```
interface interface-name {
    next-hop address;
}
```

Hierarchy Level

```
[edit forwarding-options next-hop-group group-name]
```

Description

Specify the output interface for sending copies of packets elsewhere to be analyzed.

The commit operation fails when a next-hop group has only one interface configured. It is implicitly assumed that a subgroup is up only if more than one interface in the subgroup is up.

Options

interface-name—Name of the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Next-Hop Groups to Use Multiple Interfaces to Forward Packets Used in Port Mirroring | 42

interface (Port Mirroring)

IN THIS SECTION

- Syntax | 241
- Hierarchy Level | 242
- Description | 242
- Options | 242
- Required Privilege Level | 242
- Release Information | 242

Syntax

```
interface interface-name {
    next-hop address;
}
```

Hierarchy Level

```
[edit forwarding-options port-mirroring output],
[edit forwarding-options port-mirroring family (inet | inet6) output]
```

Description

Specify the output interface for sending copies of packets elsewhere to be analyzed.

Options

interface-name—Name of the interface.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Port Mirroring | 37

l2tp-tunnel-session-identifier

IN THIS SECTION

- Syntax | 243
- Hierarchy Level | 243
- Description | 243
- Required Privilege Level | 244
- Release Information | 244

Syntax

```
l2tp-tunnel-session-identifier
```

Hierarchy Level

```
[edit forwarding-options hash-key family inet
```

Description

For better distribution in load balancing for L2TP tunneled traffic, enable `l2tp-tunnel-session-identifier` to have the hash calculation of IPv4 packets include L2TP header parameters (tunnel ID and session ID). With this option enabled, Junos OS generates different hashes for packets from different tunnels. It will also generate different hashes for packets that belong to the same tunnel, but different sessions. L2TP control traffic is not load balanced based on tunnel ID and session ID.

When the `l2tp-tunnel-session-identifier` option is not enabled, the same hash value is computed for L2TP tunneled traffic using the outer IP header, regardless of tunnel (for tunnels created between the same end points).

To help diagnosis load balancing issues, you can run the `request pfe execute command "show jnh lb" target` *target* command from the Junos OS command line.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.2.

RELATED DOCUMENTATION

hash-key (Forwarding Options) | 208

Understanding Per-Packet Load Balancing | 75

Configuring Per-Packet Load Balancing | 77

link-layer-broadcast-inet-check

IN THIS SECTION

- Syntax | 245
- Hierarchy Level | 245
- Description | 245
- Required Privilege Level | 245
- Release Information | 245

Syntax

```
link-layer-broadcast-inet-check;
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Enable destination MAC and IP address check to prevent the Router from forwarding IPV4 packets, which have link layer destination address set to `broadcast` or `multicast`, unless it is directed to an *IPV4 multicast* address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

load-balance (Forwarding Options)

IN THIS SECTION

- Syntax | 246
- Hierarchy Level | 246
- Description | 247
- Options | 248
- Required Privilege Level | 248
- Release Information | 248

Syntax

```
load-balance {
    indexed-load-balance;
    per-flow {
        hash-seed;
    }
    per-prefix {
        hash-seed number;
    }
}
```

Hierarchy Level

```
[edit forwarding-options],
[edit logical-systems logical-system-name forwarding-options],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options],
[edit routing-instances routing-instance-name forwarding-options]
```

Description

Enable per-prefix or per-flow load balancing so that the router or switch elects a next hop independently of the route selected by other routers or switches.

For the active route, when there are multiple equal-cost paths to the same destination, by default, Junos OS chooses in a random fashion one of the next-hop addresses to install into the forwarding table. Whenever the set of next hops for a destination changes in any way, the next-hop address is chosen again, also in a random fashion.

You can configure Junos OS so that, for the active route, all next-hop addresses for a destination are installed in the forwarding table. This is called per-packet load balancing. You can use load balancing to spread traffic across multiple paths between routing devices. The behavior of the per-packet load-balancing function varies according to the version of the Internet Processor ASIC in the routing device.

On routing devices with an Internet Processor I ASIC, when per-packet load balancing is configured, traffic between routing devices with multiple paths is spread in a random fashion across the available interfaces. The forwarding table balances the traffic headed to a destination, transmitting packets in round-robin fashion among the multiple next hops (up to a maximum of eight equal-cost load-balanced paths). The traffic is load-balanced on a per-packet basis.

Per-packet load distribution uses a hashing algorithm that distributes packets over equal-cost links. The algorithm is designed to distribute packets to prevent any single link from being saturated. However, per-packet load balancing offers no guarantee of equal distribution of traffic over equal-cost links, nor does it guarantee that increasing the number of Internet flows creates a better hash distribution.

On routing devices with the Internet Processor II ASIC and T Series Internet Processor II ASIC, when per-packet load balancing is configured, traffic between routing devices with multiple paths is divided into individual traffic flows (up to a maximum of 16 equal-cost load-balanced paths). On some platforms, you can increase the number of paths by using the `chassis maximum-ecmp` statement. With this statement, you can change the maximum number of equal-cost load-balanced paths to 32 or 64. Packets for each individual flow are kept on a single interface. To recognize individual flows in the transit traffic, the routing device examines each of the following:

- Source IP address
- Destination IP address
- Protocol
- Source port number
- Destination port number
- Source interface index

- Type of service (ToS)

The routing device recognizes packets in which all of these parameters are identical, and it ensures that these packets are sent out through the same interface. This prevents problems that might otherwise occur with packets arriving at their destination out of their original sequence.

Load balancing is not supported on management and internal Ethernet (**fxo**) interfaces because this type of interface cannot handle the routing process. On **fxp** interfaces, you cannot configure multiple next hops and enable load balancing.

Options

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

Support for per-flow load balancing introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

Example: Load Balancing BGP Traffic

Configuring Per-Flow Load Balancing Based on Hash Values | 87

Configuring Per-Prefix Load Balancing | 85

load-balance-group

IN THIS SECTION

- Syntax | 249
- Hierarchy Level | 249
- Description | 249
- Options | 250
- Required Privilege Level | 250
- Release Information | 250

Syntax

```
load-balance-group group-name {
    next-hop-group [ group-names ];
}
```

Hierarchy Level

```
[edit firewall]
```

Description

Configure a load-balance group.

Options

group-name—Name of load-balance group.

group-names—Name of next-hop groups to include in the load-balance group set.

Required Privilege Level

firewall—To view this statement in the configuration.

firewall-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Load-Balance Groups | 60

Routing Policies, Firewall Filters, and Traffic Policers User Guide

local-bias

IN THIS SECTION

- Syntax | 251
- Hierarchy Level | 251
- Description | 251
- Required Privilege Level | 251
- Release Information | 252

Syntax

```
local-bias percent bias;
```

Hierarchy Level

```
[edit interfaces rlt x logical-tunnel-options load-balance]
```

Description

Next hop addresses may be local or remote, and absent any other load balancing scheme (except adaptive load balancing), traffic can be expected to be more-or-less evenly distributed among the available next-hop addresses whether they are local or remote. You can skew distribution to favor local addresses, however, by setting a value for local bias (local relative to the packet forwarding engine (PFE) performing the packet look up).

For example, a value of 100 would exclude remote next-hop addresses from the traffic distribution by forcing 100% of next-hop traffic flows to use local addresses. A value of 50, on the other hand, would skew 50% of the flows that would otherwise use remote links so they use local links instead. That is, for a value set to 50, given four next-hop links, two of which are local and two of which are remote, each of the remote links could be expected to get one eighth of the flows (25% / 2) = 12.5%. Likewise, each of the local links could also be expected to receive about a third of the flows (25% + 12.5%) = 37.5%.

In contrast, with no value set for local bias, each of the four links would be expected to receive 25% of the total flows.

Required Privilege Level

interface - To view statement in the configuration.

interface-control - To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.2R1.

local-dump

IN THIS SECTION

- Syntax | 252
- Hierarchy Level | 252
- Description | 252
- Options | 253
- Required Privilege Level | 253
- Release Information | 253

Syntax

```
(local-dump | no-local-dump);
```

Hierarchy Level

```
[edit forwarding-options sampling family (inet | inet6 | mpls) output flow-server hostname]
```

Description

Enable collection of cflowd records in a log file.

Options

no-local-dump—Do not dump cflowd records to a log file before exporting.

local-dump—Dump cflowd records to a log file before exporting.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Directing Traffic Sampling Output to a Server Running the cflowd Application | 12

max-packets-per-second

IN THIS SECTION

- Syntax | 254
- Hierarchy Level | 254
- Description | 254
- Options | 254
- Required Privilege Level | 254
- Release Information | 254

Syntax

```
max-packets-per-second number;
```

Hierarchy Level

```
[edit forwarding-options sampling input]
```

Description

Specify that the traffic threshold that must be exceeded before packets are dropped. A value of 0 instructs the Packet Forwarding Engine not to sample any traffic.

Options

number—Maximum number of packets per second.

- **Range:** 0 through 65,535
- **Default:** 1000

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Traffic Sampling | 6

maximum-hop-count

IN THIS SECTION

- Syntax | 255
- Hierarchy Level | 255
- Description | 256
- Options | 256
- Required Privilege Level | 256
- Release Information | 256

Syntax

```
maximum-hop-count number;
```

Hierarchy Level

```
[edit forwarding-options helpers bootp],
[edit forwarding-options helpers bootp interface (interface-name | interface-group)]
```

Description

Set the maximum allowed number of hops. This value is compared against the hops field in the BOOTP request message. BOOTP request messages that have a number in the hops field that exceeds `maximum-hop-count` are not forwarded. If you omit the `maximum-hop-count` statement, the default value is four hops.

Options

number—Maximum number of hops for BOOTP request messages.

- **Range:** 1 through 16
- **Default:** 4

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents | **109**

maximum-packet-length

IN THIS SECTION

- Syntax | 257
- Hierarchy Level | 257
- Description | 257
- Options | 258
- Required Privilege Level | 258
- Release Information | 258

Syntax

```
maximum-packet-length bytes;
```

Hierarchy Level

```
[edit forwarding-options analyzer analyzer-name input],
[edit forwarding-options port-mirroring input],
[edit forwarding-options port-mirroring instance instance-name input],
[edit forwarding-options sampling input],
[edit forwarding-options sampling instance instance-name input]
```

Description

Set the maximum packet length to be used for port mirroring or traffic sampling. Packets longer than the maximum are truncated. This statement cannot be used with inline flow monitoring (`[edit forwarding-options sampling instance instance-name family (inet | inet6) output inline-jflow]`.

NOTE: For MX Series routers with Modular Port Interface Concentrators (MPCs), when `maximum-packet-length` (clip length) is configured for port-mirrored packets and the mirror-destination interface is a next-hop-group, the clip length is effective only for the first member interface of the next-hop-group. The mirrored packet copy sent to the rest of the interfaces is not clipped.

In addition, native analyzer sessions (that is, the `[edit forwarding-options analyzer analyzer-name input]` hierarchy level for MX Series routers) can be configured without specifying input parameters. As such, these instances use the following input values by default: rate = 1, and maximum-packet-length = 0.

Options

bytes—Maximum length (in bytes) of the mirrored packet or the sampled packet.

Set the maximum-packet-length value to zero to disable truncation; that is, to mirror or sample the entire packet. Otherwise, Juniper recommends that you configure the packet length to be equal to, or greater than, the IP header length. For IPv4, set the maximum length to at least 20, and for IPv6, set the maximum length to at least 40.

- **Range:** 0 through 9216. For MX Series routers with MPCs, and for EX9200 switches, the range is 1 through 255 bytes.
- **Default:** 0

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

For MX Series routers except the MX 80, support at the `[edit forwarding-options analyzer analyzer-name input]` hierarchy level was introduced in Junos OS Release 14.1

RELATED DOCUMENTATION

Configuring Port Mirroring | 37

Configuring Traffic Sampling on MX, M and T Series Routers

minimum-wait-time

IN THIS SECTION

- Syntax | 259
- Hierarchy Level | 259
- Description | 260
- Options | 260
- Required Privilege Level | 260
- Release Information | 260

Syntax

```
minimum-wait-time seconds;
```

Hierarchy Level

```
[edit forwarding-options helpers bootp],
[edit forwarding-options helpers bootp interface (interface-name | interface-group)]
```

Description

To set the minimum allowed number of seconds in the secs field of the BOOTP message, include the minimum-wait-time statement. This setting configures a minimum number of seconds since the client sent its first BOOTP request. BOOTP messages that have a smaller number in the secs field than the allowed minimum are not forwarded. The default value for the minimum wait time is zero (0).

The default value for the minimum wait time is zero (0) seconds. If the minimum wait time is 0 and the secs field in the BOOTP request message is 0, the device forwards the packet.

Options

seconds Minimum wait time the BOOTP client has waited before packets are forwarded.

- **Range:** 0 to 30,000
- **Default:** 0

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents | 109

mirror-once

IN THIS SECTION

- Syntax | 261
- Hierarchy Level | 261
- Description | 261
- Required Privilege Level | 261
- Release Information | 262

Syntax

```
mirror-once;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring]
```

Description

Configure the router to mirror packets only once. This feature is useful if you configure port mirroring on both ingress and egress interfaces, which could result in the same packet being mirrored twice.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3 (MX Series routers only).

Support extended to M120 routers in Junos OS Release 9.5.

RELATED DOCUMENTATION

Configuring Port Mirroring | 37

monitoring

IN THIS SECTION

- Syntax | 262
- Hierarchy Level | 263
- Description | 263
- Required Privilege Level | 263
- Release Information | 264

Syntax

```
monitoring group-name {
    family inet {
        output {
            cflowd hostname {
                port port-number;
            }
            export-format cflowd-version-5;
```

```
            flow-active-timeout seconds;
            flow-export-destination {
                (cflowd-collector | collector-pic);
            }
            flow-inactive-timeout seconds;
            interface interface-name {
                engine-id number;
                engine-type number;
                input-interface-index number;
                output-interface-index number;
                source-address address;
            }
        }
    }
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Specify flow monitoring instance name and properties.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring Passive Flow Monitoring | 35

next-hop (Forwarding Options)

IN THIS SECTION

- Syntax | 264
- Hierarchy Level | 264
- Description | 265
- Options | 265
- Required Privilege Level | 265
- Release Information | 265

Syntax

```
next-hop address;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring family (inet | inet6 | ccc | vpls) output
interface interface-name]
```

Description

Specify the next-hop address for sending copies of packets to an analyzer.

NOTE: You need not configure the option for QFX5220-128C and QFX5220-32CD routers.

Options

address—IP address of the next-hop router.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Port Mirroring | 37

next-hop-group (Forwarding Options)

IN THIS SECTION

- Syntax | 266
- Syntax EX3400, EX4100, EX4100-F, EX4300-MP, and EX4400 Switches | 266
- Hierarchy Level | 267
- Description | 267
- Options | 267
- Required Privilege Level | 268
- Release Information | 268

Syntax

```
next-hop-group group-name {
    interface interface-name {
        next-hop address;
    }
    next-hop-subgroup subgroup-name {
        interface interface-name {
            next-hop address;
        }
    }
}
```

Syntax EX3400, EX4100, EX4100-F, EX4300-MP, and EX4400 Switches

```
next-hop-group group-name {
  group-type layer-2;
```

```
  interface interface-name;
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Specify the next-hop address for sending copies of packets to an analyzer.

(This paragraph applies only to routers.) The commit operation fails when a next-hop group has only one interface configured. It is implicitly assumed that a subgroup is up only if more than one interface in the subgroup is up.

NOTE: (This Note applies only to routers.) In Junos OS releases through Release 14.2, the `next-hop-group` statement is present in the `forwarding-options` stanza for a routing instance, but the `next-hop-group` statement is not allowed in a routing instance. In other words, in a routing instance, `[edit routing-instances routing-instance-name forwarding-options next-hop-group]` is not supported. You will get an error message if you try to commit this type of configuration. Starting in Junos OS Release 14.2, the `next-hop-group` statement is not present in `[edit routing-instances routing-instance-name forwarding-options]`.

Options

group-name — Name of next-hop group.

- Routers—Up to 30 next-hop groups are supported. Each next-hop group must have at least two next-hop addresses.
- Switches—Up to 4 next-hop groups are supported.

next-hop *address* — (Applies only to routers) IP address of the next-hop router. Each next-hop group supports up to 16 next-hop addresses. Each next-hop group must have at least two next-hop addresses.

group-type layer-2 — (Applies only to EX switches) Group type of the next-hop group. The allowed value for `group-type` is `layer-2`.

interface *interface-name* — Interface used to reach the next-hop destination.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 22.4R1 for EX3400, EX4100, EX4100-F, EX4300-MP, and EX4400 switches.

RELATED DOCUMENTATION

Configuring Next-Hop Groups to Use Multiple Interfaces to Forward Packets Used in Port Mirroring | 42

next-hop-group

IN THIS SECTION

- Syntax | 269
- Hierarchy Level | 269
- Description | 269
- Options | 270
- Required Privilege Level | 270

- Release Information | 270

Syntax

```
next-hop-group group-name{
    group-type inet6;
    interface interface-name {
        next-hop ipv6-address;
    }
    next-hop-subgroup group-name{
        interface interface-name {
            next-hop ipv6-address;
        }
    }
}
```

Hierarchy Level

```
[edit forwarding-options port-mirroring family inet6 output]
```

Description

Specify the next-hop group through which to send port-mirror traffic to an analyzer. This configuration enables multipacket port mirroring on MX Series routers with or without the use of a Tunnel PIC. It is implicitly assumed that a subgroup is up only if more than one interface in the subgroup is up.

Options

group-name—Name of the next-hop group. Up to 30 next-hop groups are supported for the router. Each next-hop group must have at least two next-hop addresses.

interface-name—Name of the interface used to reach the next-hop destination.

ipv6-address—IPv6 address of the next-hop router. Each next-hop group supports up to 16 next-hop addresses. Each next-hop subgroup can have up to 16 next-hop groups.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers

no-filter-check

IN THIS SECTION

- Syntax | 271
- Hierarchy Level | 271
- Description | 271
- Required Privilege Level | 271

- Release Information | 271

Syntax

```
no-filter-check;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring output],
[edit forwarding-options port-mirroring family (inet | inet6 | ccc | vpls) output]
```

Description

Disable filter checking on the port-mirroring interface.

This statement is required when you send port-mirrored traffic to a Tunnel Services PIC that has a filter applied to it.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Port Mirroring | 37

no-listen

IN THIS SECTION

- Syntax | 272
- Hierarchy Level | 272
- Description | 273
- Required Privilege Level | 273
- Release Information | 273

Syntax

```
no-listen;
```

Hierarchy Level

```
[edit forwarding-options helpers bootp interface (interface-name | interface-group)],
[edit forwarding-options helpers domain interface interface-name],
[edit forwarding-options helpers port port-number interface interface-name],
[edit forwarding-options helpers tftp interface interface-name]
```

Description

Disable recognition of DNS requests or stop packets from being forwarded on a logical interface, a group of logical interfaces, a router, or a switch.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

- Configuring DNS and TFTP Packet Forwarding | 112
- Configuring Port-based LAN Broadcast Packet Forwarding | 117
- Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents | 109

output (Accounting)

IN THIS SECTION

- Syntax | 274
- Hierarchy Level | 274
- Description | 274
- Required Privilege Level | 275
- Release Information | 275

Syntax

```
output {
    cflowd [ hostnames ] {
        aggregation {
            autonomous-system;
            destination-prefix;
            protocol-port;
            source-destination-prefix {
                caida-compliant;
            }
            source-prefix;
        }
        autonomous-system-type (origin | peer);
        port port-number;
        version format;
    }
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    interface interface-name {
        engine-id number;
        engine-type number;
        source-address address;
    }
}
```

Hierarchy Level

```
[edit forwarding-options accounting group-name]
```

Description

Configure cflowd, output interfaces, and flow properties.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Discard Accounting | 31

output (Forwarding Table)

IN THIS SECTION

- Syntax | 275
- Hierarchy Level | 276
- Description | 276
- Options | 276
- Required Privilege Level | 276
- Release Information | 276

Syntax

```
output filter-name;
```

Hierarchy Level

```
[edit forwarding-options family (inet | inet6 | mpls) filter],
[edit routing-instances routing-instance-name forwarding-options family (inet | inet6 | mpls)
filter]
```

Description

Configure filtering on the egress traffic of the forwarding table.

Options

filter-name—Name of the applied filter.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

RELATED DOCUMENTATION

Applying Forwarding Table Filters | 54

output (Monitoring)

IN THIS SECTION

- Syntax | 277
- Hierarchy Level | 278
- Description | 278
- Required Privilege Level | 278
- Release Information | 278

Syntax

```
output {
    cflowd hostname {
        port port-number;
    }
    export-format cflowd-version-5;
    flow-active-timeout seconds;
    flow-export-destination {
        (cflowd-collector | collector-pic);
    }
    flow-inactive-timeout seconds;
    interface interface-name {
        engine-id number;
        engine-type number;
        input-interface-index number;
        output-interface-index number;
        source-address address;
    }
}
```

Hierarchy Level

```
[edit forwarding-options monitoring group-namefamily inet]
```

Description

Configure cflowd, output interfaces, and flow properties.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Passive Flow Monitoring | 35

output (Port Mirroring)

IN THIS SECTION

Syntax | 279

- Hierarchy Level | 279
- Description | 280
- Required Privilege Level | 280
- Release Information | 280

Syntax

```
output {
    interface interface-name {
        next-hop address;
    }
    next-hop-group group-name{
        group-type inet6;
        interface interface-name {
            next-hop ipv6-address;
        }
        next-hop-subgroup group-name{
            interface interface-name {
                next-hop ipv6-address;
            }
        }
    }
    no-filter-check;
    server-profile server-profile-name;
}
```

Hierarchy Level

```
[edit forwarding-options port-mirroring family (ccc | inet | inet6 | mpls | vpls)],
[edit forwarding-options port-mirroring instance instance-name family (ccc | inet | inet6 | mpls |
vpls)]
```

Description

Configure the port mirroring destination properties.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`vpls` option introduced in Junos OS Release 9.3 for MX Series routers only; support extended to M7i, M10i, M120, and M320 routers in Junos OS Release 9.5.

`ccc` option introduced in Junos OS Release 9.6 for M120 and M320 routers only.

`server-profile` option introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers only.

`next-hop-group` option introduced for `family inet6` in Junos OS Release 14.2 for MX Series routers only.

RELATED DOCUMENTATION

Configuring Port Mirroring | 37

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

output (Sampling)

IN THIS SECTION

- Syntax | 281
- Hierarchy Level | 282
- Description | 282
- Required Privilege Level | 282
- Release Information | 283

Syntax

```
output {
    aggregate-export-interval seconds;
    flow-server hostname {
        aggregation {
            autonomous-system;
            destination-prefix;
            protocol-port;
            source-destination-prefix {
                caida-compliant;
            }
            source-prefix;
        }
        autonomous-system-type (origin | peer);
        (local-dump | no-local-dump);
        port port-number;
        source-address address;
        version format;
        version9 {
            template template-name;
        }
    }
    extension-service service-name;
    file filename filename  <disable> <files number> <stamp | no-stamp> <size bytes> <world-
```

```
readable | no-world-readable>;
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    flow-server host-name {
        aggregation;
        autonomous-system-type (origin | peer);
        (local-dump | no-local-dump);
        port number;
        source-address address;
        version (5 | 8);
        version9;
    }
    interface interface-name {
        engine-id number;
        engine-type number;
        source-address address;
    }
}
```

Hierarchy Level

```
[edit forwarding-options sampling family (inet | inet6 | mpls)]
```

Description

Configure cflowd, output files and interfaces, and flow properties. Enable the collection of traffic flows using the version 9 format.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`version9` statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Collecting Traffic Sampling Output in the Cisco Systems NetFlow Services Export Version 9 Format | 16

Collecting Traffic Sampling Output in a File | 10

per-flow

IN THIS SECTION

- Syntax | 283
- Hierarchy Level | 284
- Description | 284
- Options | 284
- Required Privilege Level | 284
- Release Information | 284

Syntax

```
per-flow {
    hash-seed;
}
```

Hierarchy Level

```
[edit forwarding-options load-balance],
[edit logical-systems logical-system-name forwarding-options load-balance],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options load-balance],
[edit routing-instances routing-instance-name forwarding-options load-balance]
```

Description

Enable per-flow load balancing based on hash values.

Options

hash-seed—Configure the hash value. Junos OS automatically chooses a value for the hashing algorithm used. You cannot configure a specific hash value for per-flow load balancing.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3 (M120, M320, and MX Series routers only).

RELATED DOCUMENTATION

Configuring Per-Flow Load Balancing Based on Hash Values | 87

load-balance (Forwarding Options) | 246

per-prefix

IN THIS SECTION

- Syntax | 285
- Hierarchy Level | 285
- Description | 285
- Options | 286
- Required Privilege Level | 286
- Release Information | 286

Syntax

```
per-prefix {
    hash-seed number;
}
```

Hierarchy Level

```
[edit forwarding-options load-balance],
[edit logical-systems logical-system-name forwarding-options load-balance],
[edit logical-systems logical-system-name routing-instances routing-instance-name forwarding-
options load-balance],
[edit routing-instances routing-instance-name forwarding-options load-balance]
```

Description

Configure the hash parameter for per-prefix load balancing.

Options

hash-seed—Per-prefix load-balancing hash function.

number—Hash value.

- **Range:** 0 through 65,534
- **Default:** 0

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.0.

RELATED DOCUMENTATION

Configuring Per-Prefix Load Balancing | 85

load-balance (Forwarding Options) | 246

port (cflowd)

IN THIS SECTION

- Syntax | 287
- Hierarchy Level | 287
- Description | 287

- Options | 287
- Required Privilege Level | 287
- Release Information | 288

Syntax

```
port port-number;
```

Hierarchy Level

```
[edit forwarding-options accounting group-name output cflowd hostname],
[edit forwarding-options monitoring group-name family inet output flow-server hostname],
[edit forwarding-options sampling family (inet | inet6 | mpls) output flow-server hostname]
```

Description

Specify the UDP port number on the cflowd host system.

Options

port-number—Any valid UDP port number on the host system.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Directing Traffic Sampling Output to a Server Running the cflowd Application | 12

port (Packet Forwarding)

IN THIS SECTION

- Syntax | 288
- Hierarchy Level | 289
- Description | 289
- Options | 290
- Required Privilege Level | 290
- Release Information | 290

Syntax

```
port port-number {
    description text-description;
    interface interface-name {
        broadcast;
        description text-description;
        no-listen;
        server address <logical-system logical-system-name> <routing-instance routing-instance-
name>;
    }
```

```
    server address <logical-system logical-system-name> <routing-instance routing-instance-name>;
}
```

Hierarchy Level

```
[edit forwarding-options helpers]
```

Description

Configure a port helper on the router or switch, which listens for LAN broadcast traffic on a custom UDP port number and forwards traffic to particular destination servers as unicast traffic.

To set up a port helper, configure the UDP port number and optionally an interface on which to listen for broadcast traffic, and the destination server address to receive that traffic, as shown in either of the following sample configurations:

```
user@ host# show forwarding-options
helpers {
    port 1200 {
        server 10.20.30.40;
        }
    }
```

```
user@ host# show forwarding-options
helpers {
    port 3000 {
        interface {
            fe-0/0/1.0 {
                server 192.0.2.2;
            }
        }
    }
    port 3001 {
        interface {
            fe-0/0/0.0 {
```

```
                server 192.0.2.2;
            }
        }
    }
}
```

Starting in Junos OS Release 17.2R1, you can configure forwarding traffic to multiple destination servers for a given port number by specifying multiple port configuration statements with the same port number and different server addresses.

You cannot configure port helpers for standard ports used by services such as BOOTP, DNS and TFTP; instead, use the `helpers` configuration statements specifically for forwarding packets for those protocols.

The remaining statements are explained separately. See CLI Explorer.

Options

port-number	UDP port number for listening. • **Range:** 1 through 65535

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for multiple server instances for a given port introduced in Junos OS Release 17.2 for MX Series routers.

Support for multiple servers on a given port introduced in Junos OS Release 17.3R1 for EX9200 switches.

RELATED DOCUMENTATION

Configuring DNS and TFTP Packet Forwarding | 112

Configuring Port-based LAN Broadcast Packet Forwarding | 117

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents | 109

port-mirroring

IN THIS SECTION

- Syntax: MX Series and PTX Series Routers, M120 and M320 | 291
- Syntax: QFX Series (but not QFX10000) and EX4600 Switches and NFX Series Devices | 293
- Syntax: OCX1100 | 294
- Syntax: QFX10000 Switches | 295
- Hierarchy Level | 295
- Description | 295
- Required Privilege Level | 295
- Release Information | 296

Syntax: MX Series and PTX Series Routers, M120 and M320

```
port-mirroring {
    input {
        maximum-packet-length bytes;
        rate number;
        run-length number;
    }
    family (any | ccc | inet | inet6 | vpls) {
        output {
            interface interface-name {
                next-hop address;
            }
            next-hop-group group-name{
```

```
                group-type inet6;
                interface interface-name {
                    next-hop ipv6-address;
                }
                next-hop-subgroup group-name{
                    interface interface-name {
                        next-hop ipv6-address;
                    }
                }
            }
            no-filter-check;
        }
    }
    instance {
        instance-name {
            input {
                maximum-packet-length bytes;
                rate number;
                run-length number;
            }
            family (any | ccc | inet | inet6 | vpls) {
                output {
                    interface interface-name {
                        next-hop address;
                    }
                    no-filter-check;
                    server-profile server-profile-name;
                }
            }
        }
    }
    mirror-once;
    traceoptions {
        file filename  <files number> <size bytes> <world-readable | no-world-readable>;
    }
}
```

Syntax: QFX Series (but not QFX10000) and EX4600 Switches and NFX Series Devices

```
port-mirroring {
    family {
        ethernet-switching
            output {
                interface interface-name {
                }
                no-filter-check;
                }
                vlan vlan-name {
                    no-tag;
                }
            }
        inet
            output {
                ip-address address {
                }
                routing-instance instance-name {
                    ip-address address {
                }
            }
    }
    instance instance-name {
        family
            ethernet-switching {
                output {
                    interface interface-name {
                    }
                    no-filter-check;
                    }
                    vlan vlan-name {
                        no-tag;
                    }
                }
            inet
                output {
                    ip-address address {
                    }
                    routing-instance instance-name {
```

```
                        ip-address address {
                    }
                }
            }
        }
    }
}
```

Syntax: OCX1100

```
port-mirroring {
    family {
        inet
            output {
                ip-address address {
                }
                routing-instance instance-name {
                    ip-address address {
                }
            }
    }
    instance instance-name {
        family
            inet
                output {
                    ip-address address {
                    }
                    routing-instance instance-name {
                        ip-address address {
                    }
                }
            }
        }
    }
}
```

Syntax: QFX10000 Switches

```
port-mirroring {
     instance instance-name {
              input {
                   rate number;
                   run-length number;
              }
              output {
                  dscp value;
                  forwarding-class class-name;
                  ip-destination-address address;
                  ip-source-address address;
                  policer policer-name;
              }
}
```

Hierarchy Level

```
[edit  forwarding-options]
```

Description

Create a port-mirroring configuration. Specify the address family, rate, run length, interface, and next-hop address for sending copies of packets to an analyzer.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`family vpls` statement introduced in Junos OS Release 9.3.

`instance` *`instance-name`* statement introduced in Junos OS Release 9.3 .

`mirror-once` statement introduced in Junos OS Release 9.3 .

`family ccc` statement introduced in Junos OS Release 9.6.

`family inet6` and `next-hop-group` statements introduced in Junos OS Release 14.2 .

RELATED DOCUMENTATION

Configuring Port Mirroring | 37

Configuring Port Mirroring

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

Understanding Port Mirroring and Analyzers

Understanding Port Mirroring

Examples: Configuring Port Mirroring for Local Analysis

Example: Mirroring Employee Web Traffic with a Firewall Filter

rate (Forwarding Options)

IN THIS SECTION

- Syntax | 297
- Hierarchy Level | 297
- Description | 297

Options | 298
Required Privilege Level | 298
Release Information | 298

Syntax

```
rate number;
```

Hierarchy Level

```
[edit forwarding-options analyzer analyzer-name input],
[edit forwarding-options port-mirroring input],
[edit forwarding-options sampling input],
[edit forwarding-options sampling instance instance-name input]
```

Description

Set the ratio of the number of packets to be sampled. For example, if you specify a rate of 10, every tenth packet (1 packet out of 10) is sampled.

Native analyzer sessions (that is, the `[edit forwarding-options analyzer analyzer-name input]` hierarchy level for MX Series routers) can be configured without specifying input parameters, which means that the instance uses default input values: rate = 1 and maximum-packet-length = 0.

NOTE: The recommended sampling rate for the MX150 is 1000 or greater. If you configure less than 1000, a warning is issued.

Options

number—Denominator of the ratio.

- **Range:** 1 through 16000000 (16M)

  For QFX Series switches, the maximum sampling rate for inline Junos Traffic Vision (J-Flow) is 65535.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the `[edit forwarding-options analyzer analyzer-name input]` hierarchy level for MX Series routers introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring Port Mirroring | 37

Configuring Traffic Sampling | 6

relay-agent-option

IN THIS SECTION

Syntax | 299

Hierarchy Level | 299

- Description | 299
- Required Privilege Level | 299
- Release Information | 300

Syntax

```
relay-agent-option;
```

Hierarchy Level

```
[edit forwarding-options helpers bootp],
[edit logical-systems routing-instances instance-name forwarding-options helpers bootp],
[edit routing-instances instance-name forwarding-options helpers bootp]
```

Description

Enable the DHCP relay agent information option which allows DHCP to forward information from clients on different VRF routing instances. The functionality is described in RFC 3046, *DHCP Relay Agent Information Option*. For the Junos OS implementation, the DHCP option number is 82, and the suboption ID is 1. The suboption length is the length required to contain an interface name in addition to the terminating null character. The overall option length is the suboption length plus 2 bytes (for the option header). The DHCP relay agent information option is only present on packets sent between the relay and the DHCP server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents | 109

route-accounting

IN THIS SECTION

- Syntax | 300
- Hierarchy Level | 300
- Description | 301
- Default | 301
- Required Privilege Level | 301
- Release Information | 302

Syntax

```
route-accounting;
```

Hierarchy Level

```
[edit forwarding-options family inet6]
```

Description

Enables the collection of transit IPv6 statistics at both the physical interface and logical interface levels on third-generation FPCs (FPC3-PTX-U2 and FPC3-PTX-U3 on PTX5000 and FPC3-SFF-PTX-U0 and FPC3-SFF-PTX-U1 on PTX3000), PTX1000, and PTX10008. Enable this statement prior to the route being created.

Egress accounting for IPv6 traffic is not performed for cases where MPLS packets arrive on a TCC interface and then egress the router as IPv6 packets.

PTX Series routers with first or second-generation FPCs do not display IPv6statistics for physical interfaces or logical interfaces. Transit statistics on child links in aggregated Ethernet interfaces are not included.

To view the IPv6 statistics for a given physical or logical interface, run the operational command, `show interfaces extensive interface | find IPv6` or `show interfaces statistics` *`interface-name`* `detail`. Note that the output displays packet and byte counts for transit traffic only. Locally generated packets are not included in the metrics.

NOTE:

- Enabling this option soon after disabling it on a third-generation FPCs (FPC3-PTX-U2 and FPC3-PTX-U3 on PTX5000 and FPC3-SFF-PTX-U0 and FPC3-SFF-PTX-U1 on PTX3000), PTX1000, and PTX10008, restarts the accounting from zero and does not retain the statistics that was previously accounted.
- On third-generation FPCs in PTX series, egress accounting for IPV6 traffic is not performed for cases where MPLS packets arrives on TCC interface and egress out of the router as IPV6 packets.

Default

Disabled

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Configuring IPv4 and IPv6 Accounting | 29

run-length

IN THIS SECTION

- Syntax | 302
- Hierarchy Level | 302
- Description | 303
- Options | 303
- Required Privilege Level | 303
- Release Information | 303

Syntax

```
run-length number;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring input],
[edit forwarding-options port-mirroring instance port-mirroring-instance-name input],
[edit forwarding-options  port-mirroring family (inet|inet6) input],
```

```
[edit forwarding-options sampling input],
[edit forwarding-options sampling instance instance-name input]
```

Description

Set the number of samples following the initial trigger event. The configuration enables you to sample packets following those already being sampled.

NOTE: The `run-length` statement is not supported when you configure inline flow monitoring (by including the `inline-jflow` statement at the `[edit forwarding-options sampling instance instance-name family (inet | inet6) output]` hierarchy level).

Options

number—Number of samples.

- **Range:** 0 through 20
- **Default:** 0

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying Forwarding Table Filters | 54

Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers

Configuring Traffic Sampling on MX, M and T Series Routers

sampling (Forwarding Options)

IN THIS SECTION

- Syntax | 304
- Hierarchy Level | 307
- Description | 307
- Required Privilege Level | 307
- Release Information | 307

Syntax

```
sampling {
    disable;
    family (inet | inet6 | mpls | vpls) {
        disable;
        output {
            aggregate-export-interval seconds;
            extension-service service-name;
            file {
                disable;
                filename filename;
                files number;
                size bytes;
                (stamp | no-stamp);
                (world-readable | no-world-readable);
            }
            flow-active-timeout seconds;
```

```
            flow-inactive-timeout seconds;
            flow-server hostname {
                aggregation {
                    autonomous-system;
                    destination-prefix;
                    protocol-port;
                    source-destination-prefix {
                        caida-compliant;
                    }
                    source-prefix;
                }
                autonomous-system-type (origin | peer);
                (local-dump | no-local-dump);
                port port-number;
                source-address address;
                version format;
                version9 {
                    template template-name;
                }
            }
            interface interface-name {
                engine-id number;
                engine-type number;
                source-address address;
            }
        }
    }
    input {
        max-packets-per-second number;
        maximum-packet-length bytes;
        rate number;
        run-length number;
    }
    instance instance-name {
        disable;
        family (bridge | inet | inet6 | mpls | vpls) {
            disable;
            output {
                aggregate-export-interval seconds;
                extension-service service-name;
                flow-server hostname {
                    aggregation {
                        autonomous-system;
```

```
                        destination-prefix;
                        protocol-port;
                        source-destination-prefix {
                            caida-compliant;
                        }
                        source-prefix;
                    }
                    autonomous-system-type (origin | peer);
                    dscp dscp-value;
                    forwarding-class class-name;
                    (local-dump | no-local-dump);
                    port port-number;
                    source-address address;
                    version format;
                    version9 {
                        template template-name;
                    }
                    version-ipfix {
                        template template-name;
                    }
                }
                inline-jflow {
                    source-address address;
                    flow-export-rate rate;
                }
                interface interface-name {
                    engine-id number;
                    engine-type number;
                    source-address address;
                }
            }
        }
        input {
            max-packets-per-second number;
            maximum-packet-length bytes;
            rate number;
            run-length number;
        }
    }
    pre-rewrite-tos;
    sample-once;
    traceoptions {
        no-remote-trace;
```

```
        file filename <files number> <size bytes> <match expression> <world-readable | no-world-
readable>;
    }
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Configure traffic sampling.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support for IPv6 addresses for the `flow-server` *address* and `source-address` *address* statements within sampling instances introduced in Junos OS Evolved Release 23.1R1.

RELATED DOCUMENTATION

Configuring Traffic Sampling on MX, M and T Series Routers

Applying Forwarding Table Filters | 54

Collecting Traffic Sampling Output in the Cisco Systems NetFlow Services Export Version 9 Format | 16

Directing Traffic Sampling Output to a Server Running the cflowd Application | 12

Configuring Port Mirroring | 37

Tracing Traffic-Sampling Operations | 22

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

server (DHCP and BOOTP Relay Agent)

IN THIS SECTION

- Syntax | 308
- Hierarchy Level | 309
- Description | 309
- Options | 309
- Required Privilege Level | 309
- Release Information | 309

Syntax

```
server address {
    logical-system logical-system-name <routing-instance [ <default> routing-instance-names ]>;
    routing-instance [ <default> routing-instance-names ];
}
```

Hierarchy Level

```
[edit forwarding-options helpers bootp],
[edit forwarding-options helpers bootp interface (interface-name | interface-group)]
```

Description

Configure the router or switch to act as a DHCP and BOOTP relay agent. The device forwards all broacast requests within the configured subnet to all configured servers in parallel. To support clients on different VRFs, see the "relay-agent-option" on page 298 statement.

Options

- *address*—One or more addresses of the server.
- logical-system *logical-system-name*—(Optional) Logical system of the server.
- routing-instance *routing-instance-names*—(Optional) Routing instance name that belong to the DHCP or BOOTP relay agent.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents | 109

relay-agent-option | 298

server (DNS, Port, and TFTP Service)

IN THIS SECTION

- Syntax | 310
- Hierarchy Level | 310
- Description | 311
- Options | 311
- Required Privilege Level | 311
- Release Information | 311

Syntax

```
server address <logical-system logical-system-name> <routing-instance routing-instance-name>;
```

Hierarchy Level

```
[edit forwarding-options helpers domain],
[edit forwarding-options helpers domain interface interface-name],
[edit forwarding-options helpers port port-number],
[edit forwarding-options helpers port port-number interface interface-name],
[edit forwarding-options helpers tftp],
[edit forwarding-options helpers tftp interface interface-name]
```

Description

Specify the DNS or TFTP server for forwarding DNS or TFTP requests, or specify a destination server address for forwarding LAN broadcast packets as unicast traffic for a custom-configured UDP port.

When configuring port helpers, in releases prior to Junos OS Release 17.2, only one server can be specified for a given port. For Junos OS Release 17.2 and later, multiple servers can be specified for a given port at the global or interface-specific level. When multiple servers are specified, the same packet, with the originator IP address and port requests, is forwarded to the different configured servers; the payload of the UDP packet is not modified.

Options

address—IP address of the server.

logical-system ***logical-system-name***—(Optional) Logical system name of the server.

routing-instance [***routing-instance-names*** **]**—(Optional) Set the routing instance name or names that belong to the DNS server or TFTP server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced for port helpers in Junos OS Release 17.2R1 for EX4300 switches.

Support for multiple server instances for a given port introduced in Junos OS Release 17.2 for MX Series routers.

Support for multiple server instances for a given port introduced in Junos OS Release 17.3R1 for EX9200 switches.

RELATED DOCUMENTATION

Configuring DNS and TFTP Packet Forwarding | 112

Configuring Port-based LAN Broadcast Packet Forwarding | 117

server-address (Hosted Services)

IN THIS SECTION

- Syntax | 312
- Hierarchy Level | 312
- Description | 312
- Options | 313
- Required Privilege Level | 313
- Release Information | 313

Syntax

```
server-address ipv4-address;
```

Hierarchy Level

```
[edit services hosted-services server-profile server-profile-name]
```

Description

Configure the server address where sampled packets are sent.

Options

ipv4-address—IPv4 address of the server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

server-profile

IN THIS SECTION

- Syntax | 314
- Hierarchy Level | 314
- Description | 314
- Options | 314
- Required Privilege Level | 314
- Release Information | 314

Syntax

```
server-profile server-profile-name {
    client-address ipv4-address;
    server-address ipv4-address;
}
```

Hierarchy Level

```
[edit services hosted-services]
```

Description

Configure the server profile.

Options

server-profile-name—Name to apply to this server profile.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

server-profile (Active Flow Monitoring)

IN THIS SECTION

- Syntax | 315
- Hierarchy Level | 315
- Description | 315
- Options | 316
- Required Privilege Level | 316

Syntax

```
server-profile server-profile-name;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring instance instance-name family (inet | inet6 | mpls)
output]
```

Description

Specify the name of a server profile. This profile specifies a host where sampled traffic is sent.

Options

server-profile-name Specify the name of a server profile configured at the [`edit services hosted-services server-profile` *`server-profile-name`*] hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

hosted-services | 218

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

size (Sampling and Traceoptions)

IN THIS SECTION

- Syntax | 317
- Hierarchy Level | 317
- Description | 317
- Options | 317
- Required Privilege Level | 318
- Release Information | 318

Syntax

```
size bytes;
```

Hierarchy Level

```
[edit forwarding-options helperstraceoptions file],
[edit forwarding-options port-mirroring traceoptions file],
[edit forwarding-options sampling family family-name output file],
[edit forwarding-options sampling traceoptions file]
```

Description

Specify the maximum size of each file containing sample or log data. The file size is limited by the number of files to be created and the available hard disk space.

When a traffic sampling file named **sampling-file** reaches the maximum size, it is renamed **sampling-file.0**. When the **sampling-file** file again reaches its maximum size, **sampling-file.0** is renamed **sampling-file.1** and **sampling-file** is renamed **sampling-file.0**. This renaming scheme continues until the maximum number of traffic sampling files is reached. Then the oldest traffic sampling file is overwritten.

Options

bytes—Maximum size of each traffic sampling file or trace log file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).

- **Syntax:** ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB
- **Range:** 10 KB through the maximum file size supported on your router
- **Default:** 1 MB for sampling data; 128 KB for log information

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Collecting Traffic Sampling Output in a File | 10

Tracing Traffic-Sampling Operations | 22

source-checking

IN THIS SECTION

- Syntax | 318
- Hierarchy Level | 319
- Description | 319
- Required Privilege Level | 319
- Release Information | 319

Syntax

```
source-checking;
```

Hierarchy Level

```
[edit forwarding-options family inet6]
```

Description

(MX Series 5G Universal Routing Platforms Only) Discard IPv6 packets when the source address type is unspecified, loopback, multicast, or link-local unicast.

RFC 4291, *IP Version 6 Addressing Architecture*, refers to four address types that require special treatment when they are used as source addresses. The four address types are:

- Unspecified--The IPv6 unspecified address is 0:0:0:0:0:0:0:0, which can be notated as ::/128.
- Loopack--The IPv6 loopback address is 0:0:0:0:0:0:0:1, which can be notated as ::1/128.
- Multicast--IPv6 multicast addresses use the prefix FF00::/8.
- Link-Local--IPv6 link-local addresses have the prefix FE80::/10.

The loopback and multicast addresses must never be used as a source address in IPv6 packets. The unspecified and link-local addresses can be used as source addresses but routers must never forward packets that have these addresses as source addresses. Typically, packets that contain unspecified or link-local addresses as source addresses are delivered to the local host. If the destination is not the local host, then the packet must not be forwarded. Configuring this statement filters or discards IPv6 packets of these four address types.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

Applying Forwarding Table Filters | 54

stamp

IN THIS SECTION

- Syntax | 320
- Hierarchy Level | 320
- Description | 320
- Default | 321
- Options | 321
- Required Privilege Level | 321
- Release Information | 321

Syntax

```
(stamp | no-stamp);
```

Hierarchy Level

```
[edit forwarding-options sampling family family-name output file]
```

Description

Include a timestamp with each line in the output file.

Default

no-stamp

Options

no-stamp—Do not include timestamps.

stamp—Include a timestamp with each line of packet sampling information.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Collecting Traffic Sampling Output in a File | 10

tftp

IN THIS SECTION

- Syntax | 322
- Hierarchy Level | 322

- Description | 322
- Required Privilege Level | 323
- Release Information | 323

Syntax

```
tftp {
    description text-description;
    interface interface-name {
        broadcast;
        description text-description;
        no-listen;
        server address <logical-system logical-system-name> <routing-instance routing-instance-
name>;
    }
    server address <logical-system logical-system-name> <routing-instance routing-instance-name>;
}
```

Hierarchy Level

```
[edit forwarding-options helpers]
```

Description

Enable TFTP request packet forwarding.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring DNS and TFTP Packet Forwarding | 112

traceoptions (DNS, Port, and TFTP Packet Forwarding)

IN THIS SECTION

- Syntax | 324
- Hierarchy Level | 324
- Description | 324
- Default | 324
- Options | 324
- Required Privilege Level | 326
- Release Information | 326

Syntax

```
traceoptions {
    file filename  <files number> <match regular-expression> <size bytes> <world-readable | no-
world-readable>;
    flag flag;
    level level;
    <no-remote-trace>;
}
```

Hierarchy Level

```
[edit forwarding-options helpers]
```

Description

Configure tracing operations for BOOTP, DNS, TFTP, or custom UDP port packet forwarding.

Default

If you do not include this statement, no tracing operations are performed.

Options

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks (" "). All files are placed in a file named **fud** in the directory **/var/log**. If you include the `file` statement, you must specify a filename.

files ***number***—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

- **Range:** 2 through 1000
- **Default:** 3 files

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple `flag` statements. You can include the following flags:

- **address**—Trace address management events
- **all**—Trace all events
- **bootp**—Trace BOOTP or DHCP services events
- **config**—Trace configuration events
- **domain**—Trace DNS service events
- **ifdb**—Trace interface database operations
- **io**—Trace I/O operations
- **main**—Trace main loop events
- **port**—Trace arbitrary protocol events
- **rtsock**—Trace routing socket operations
- **tftp**—Trace TFTP service events
- **trace**—Trace tracing operations
- **ui**—Trace user interface operations
- **util**—Trace miscellaneous utility operations

match ***regular-expression***—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—(Optional) Disable remote tracing globally or for a specific tracing operation.

no-world-readable—(Optional) Restrict file access to the owner.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** file again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and filename.

- **Syntax:** *x***k** to specify KB, *x***m** to specify MB, or *x***g** to specify GB
- **Range:** 0 bytes through 4,294,967,295 KB
- **Default:** 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement standardized and **match** option introduced in Junos OS Release 8.0.

RELATED DOCUMENTATION

Configuring DNS and TFTP Packet Forwarding | 112

traceoptions (Port Mirroring and Traffic Sampling)

IN THIS SECTION

- Syntax | 327
- Hierarchy Level | 327
- Description | 327

Required Privilege Level | 327
Release Information | 328

Syntax

```
traceoptions {
    file filename  <files number> <size bytes> <world-readable | no-world-readable>;
}
```

Hierarchy Level

```
[edit forwarding-options port-mirroring],
[edit forwarding-options sampling]
```

Description

Configure traffic sampling tracing operations.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Tracing Traffic-Sampling Operations | 22

vector routing

vector-routing

IN THIS SECTION

- Syntax | 328
- Junos OS Hierarchy Level | 329
- Description | 329
- Options | 330
- Required Privilege Level | 332
- Release Information | 332

Syntax

```
vector-routing;
```

Junos OS Hierarchy Level

```
[edit services vector-routing]
```

Description

The Service Vector Routing (SVR) control plane is a location independent, service-based data model that provides the language used to describe the network's services, tenancy, and associated policies. Every router in an SVR fabric shares the same service-based policies and topology. As such, to interoperate with SVR, you need to configure the connected SRX Series Services Gateway or the NFX Series Network Services Platform so it uses the same set of device names and IP numbers so the devices can recognize each other. Do that in Junos, at the `service vector-routing` hierarchy level.

Key concepts in the data model are explained below, followed by the Junos CLI hierarchy where they appear.

- The Authority represents the data model used to describe network (routers), service, and policy behavior organized into a configuration hierarchy. In short, it represents the collection of all SVR networking components.
- Global data applies to all devices within the authority.
- Local data applies at the level of router (it is not used beyond the device).
- A Node is essentially a running instance of the SVR control plane, either physical or virtual, at the local level.
- A Peer is an SVR participating router that is connected to the one you are configuring. (You can use the notion of Peers to help diagram the network of SVR nodes.)
- Routers in this context are logical entities; they can be standalone (one node) or high-availability (two node), and are responsible for receiving and sending packets to their correct destinations. Any options configured under the router hierarchy are considered local data (applicable to that router only).
- Source Tenants , or Tenants are used in the data model as a way to place traffic sources and routes and their services (also referred to as service routes) into logical partitions within the underlay network. In short, Tenants are a way to group devices and create a trust zone.
- Device Interface, also known as Physical Interface, and found at the **interface** level of the Junos hierarchy.

- Network Interface, also known as Logical Interface, and found at the **interface** level of the Junos hierarchy.
- Adjacency is one of several ways you can tell the router how to reach its Peer (because peers can be reachable from multiple network interfaces).
- Service, in the SVR data model, is a traffic destination being accessed by constituents in your network. It is global data. Services represent specific applications that a network delivers, such as web services, database services, or voice/video services, and for which you can specify delivery limits such as latency, packet loss, and jitter.
- Service Routes are local data that is used to tell the router how to reach a particular service, like a service oriented routing table. The destination can be one or more peers, a gateway, an IP address, a subnet, or something else. However, when the destination is another SVR router, the traffic will use SVR to send the traffic.
- BFD, in the SVR network, devices use a lightweight Bi-directional Forwarding Detection (BDF) metadata to evaluate link quality and service changes between waypoints (you can think of waypoints as the IP addresses, or interfaces, used to forward SVR traffic along the SVR path).
- The 128T concepts of Service Routers Policy and Service Policy are not currently used in Junos for NFX or SRX devices.
- Neighborhoods this 128T concept is not currently used in Junos for NFX or SRX devices.
- Vector this 128T concept is not currently used in Junos for NFX or SRX devices.

Options

`authority-name` *`name`*—Specify the name of the authority the device belongs to.

- `cipher-suites` *`cipher-suite name`*—The cipher-suite you choose must be supported by both ends of the connection (the cipher suite itself is a set of cryptographic algorithms that are used by the devices)
  - authentication-algorithm
  - authentication-disabled
  - authentication-key
  - encryption-algorithm
  - encryption-disabled
  - encryption-key

- encryption-vector

`destination-services` *`name`*—Identify the destination service by giving it a name and specifying an IP prefix, transport protocol, and port.

- `access-policy` *`source-tenant-name`*—Identify the source-tenant from which the device will receive traffic.
  - permission *allow | deny*
- `ip-prefix` *`value`*—The address prefix of the destination-service.
- `transport` *`gre | icmp | tcp | udp`*—Specify the traffic protocol.

`meta-bfd` —Configure BFD to evaluate connection and path attributes between waypoints in a SVR session path.

- `desired-tx-interval` *`milliseconds`*—Specify the frequency at which to send the BFD asynchronous control packets.
- `dscp` *`value`*—Specify the DSCP value to use with BFD packets, from zero to 63.
- `link-test-interval` *`seconds`*—Specify the interval between BFD echo tests sent to the peer.
- `multiplier` *`value`*—Number of consecutive undelivered or unacknowledged meta-bdf-packets needed to qualify the link as unusable, from 3 to 20.
- `required-min-rx-interval` *`milliseconds`*—Specify the lower threshold at which to send BFD control packets.

`router` *`name`*—Name of the router used for vector-routing (explained above).

- `node` *`name`*—Name of the service node (explained above).
  - `adjacency` *`name`*—Identify adjacent routers that are reachable from this node.
    - `address` *`value`*—The IP address or hostname of the adjacent router, or the waypoint address of the peer router.
    - `cipher-suite` *`cipher-suite name`*—The cipher-suite must be supported by both ends of the connection.
    - `cost` *`value`*—You can assign an administrative cost to the link, for example to influence path preference as part of a service level agreement. The default is 0.
    - `peer` *`name`*— Remote router.
    - `peer-connectivity` *`bidirectional | outbound-only`*
- `peer` *`name`*—Name of the peer router (explained above) to which the waypoint address belongs.

- `peer-connectivity` *`bidirectional | outbound-only`*—Set to *outbound-only* if the adjacency is behind a NAT device.
- `service-route` *`name`*—Name of the target route for handling traffic for a given service.

`source-tenant` *`name`*—Name of the tenant to which you want to map traffic (using a given IP prefix or by specifying the interface the traffic will arrive from). If a given service does not match any tenants, they are considered "tenantless," which means no service policy other than public-level one will be applied to the traffic.

- `ip-prefixs` *`value`*—Specify the IP prefix(es) you want to associate with the tenant name, and thus, by extension, the traffic flows you want to apply service policy to.
- `interface` *`value`*—Specify the interface you want to associate with the tenant name.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

SVR support introduced in Junos OS Release 21.4R1 for SRX and NFS devices.

RELATED DOCUMENTATION

> https://docs.128technology.com/docs/about_128t
>
> https://docs.128technology.com/docs/config_reference_guide/

8

CHAPTER

traceoptions (Port Mirroring and Traffic Sampling)

transport-type | 334

version | 335

version9 | 337

world-readable (Forwarding Options) | 338

transport-type

IN THIS SECTION

- Syntax | 334
- Hierarchy Level | 334
- Description | 334
- Options | 334
- Required Privilege Level | 335
- Release Information | 335

Syntax

```
transport-type type;
```

Hierarchy Level

```
[edit services hosted-services server-profile server-profile-name]
```

Description

Configure the transport type.

Options

type—Transport type.

- **Range:** GRE, TCP, or UDP

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

version

IN THIS SECTION

- Syntax | 336
- Hierarchy Level | 336
- Description | 336
- Options | 336
- Required Privilege Level | 336
- Release Information | 336

Syntax

```
version format;
```

Hierarchy Level

```
[edit forwarding-options accounting group-name output cflowd hostname],
[edit forwarding-options sampling family (inet | inet6 | mpls) output flow-server hostname]
```

Description

Specify the version format of the aggregated flows exported to a cflowd server.

Options

format—Export format of the flows.

- **Values:** 5 or 8
- **Default:** 5

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Directing Traffic Sampling Output to a Server Running the cflowd Application | 12

version9

IN THIS SECTION

- Syntax | 337
- Hierarchy Level | 337
- Description | 337
- Options | 338
- Required Privilege Level | 338
- Release Information | 338

Syntax

```
version9 {
    template template-name;
}
```

Hierarchy Level

```
[edit forwarding-options sampling family family-name output flow-server hostname]
```

Description

Enable active flow monitoring using the version 9 template format to collect traffic flows.

Options

template ***template-name***—Name of a version 9 record flow format template configured at the `[edit services monitoring]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-level—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

Collecting Traffic Sampling Output in the Cisco Systems NetFlow Services Export Version 9 Format | 16

Monitoring, Sampling, and Collection Services Interfaces User Guide

Junos OS Services Interfaces Library for Routing Devices

world-readable (Forwarding Options)

IN THIS SECTION

- Syntax | 339
- Hierarchy Level | 339
- Description | 339
- Default | 339
- Options | 339

Required Privilege Level | 340
Release Information | 340

Syntax

```
(world-readable | no-world-readable);
```

Hierarchy Level

```
[edit forwarding-options helpers traceoptions file],
[edit forwarding-options port-mirroring traceoptions file],
[edit forwarding-options sampling family family-name output file],
[edit forwarding-options sampling traceoptions file]
```

Description

Enable unrestricted file access.

Default

no-world-readable

Options

no-world-readable—Restrict file access to the owner.

world-readable—Enable unrestricted file access.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Collecting Traffic Sampling Output in a File | 10

Tracing Traffic-Sampling Operations | 22

CHAPTER 9

Operational Commands

clear passive-monitoring statistics | 343

clear services flow-collector statistics | 344

request services flow-collector change-destination primary interface | 346

request services flow-collector change-destination secondary interface | 348

request services flow-collector test-file-transfer | 350

show chassis forwarding | 352

show forwarding-options hyper-mode | 355

show forwarding-options load-balance | 357

show forwarding-options port-mirroring | 362

show forwarding-options next-hop-group | 365

show interfaces (Flow Monitoring) | 370

show interfaces (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet) | 378

show interfaces statistics | 403

show passive-monitoring error | 422

show passive-monitoring flow | 426

show passive-monitoring memory | 429

show passive-monitoring status | 432

show passive-monitoring usage | 435

show route forwarding-table | 438

show services accounting aggregation | 453

show services accounting aggregation template | 459

show services accounting errors | 462

show services accounting flow | 469

show services accounting flow-detail | 478

show services accounting memory | 485

show services accounting packet-size-distribution | 488

show services accounting status | 491

show services accounting usage | 497

show services flow-collector file interface | 500

show services flow-collector input interface | 504

show services flow-collector interface | 507

clear passive-monitoring statistics

IN THIS SECTION

- Syntax | 343
- Description | 343
- Options | 343
- Required Privilege Level | 344
- Output Fields | 344
- Sample Output | 344
- Release Information | 344

Syntax

```
clear passive-monitoring statistics (all | interface interface-name)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Clear statistics for one passive monitoring interface or for all passive monitoring interfaces.

Options

all	Clear statistics for all configured passive monitoring interfaces.
interface *interface-name*	Clear statistics for the specified passive monitoring interface (**mo-***fpc/pic/port*).

Required Privilege Level

network

Output Fields

When you enter this command, you are not provided feedback on the status of your request.

Sample Output

clear passive-monitoring statistics

```
user@host> clear passive-monitoring statistics interface mo-5/0/0
```

Release Information

Command introduced in Junos OS Release 7.6.

clear services flow-collector statistics

IN THIS SECTION

- Syntax | 345
- Description | 345
- Options | 345
- Required Privilege Level | 345
- Output Fields | 345
- Sample Output | 346

- Release Information | 346

Syntax

```
clear services flow-collector statistics (all | interface interface-name)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Clear statistics for one flow collector interface or for all flow collector interfaces.

Options

all	Clear statistics for all configured flow collector interfaces.
interface *interface-name*	Clear statistics for the specified flow collector interface (**cp-***fpc/pic/port*).

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services flow-collector statistics

```
user@host> clear services flow-collector statistics interface cp-5/0/0
Flow collector interface: cp-5/0/0
Interface state: Collecting flows
Statistics cleared successfully
```

Release Information

Command introduced before Junos OS Release 7.4.

request services flow-collector change-destination primary interface

IN THIS SECTION

- Syntax | 347
- Description | 347
- Options | 347
- Required Privilege Level | 347
- Output Fields | 348
- Sample Output | 348
- Release Information | 348

Syntax

```
request services flow-collector change-destination primary interface cp-fpc/pic/port
<clear-files>
<clear-logs>
<immediately | gracefully>
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Switch to the primary File Transfer Protocol (FTP) server that is configured as a flow collector.

Options

none	Switch to the primary FTP server.
cp-*fpc/pic/port*	Use the specified flow collector interface name for the primary destination.
clear-files	(Optional) Request clearing of existing data files in the FTP wait queue when the switch takes place.
clear-logs	(Optional) Request clearing of existing logs when the switch takes place.
immediately | gracefully	(Optional) Specify whether you want the switch to take place immediately, or to affect only newly created files.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services flow-collector change-destination primary interface

```
user@host> request services flow-collector change-destination primary interface
cp-6/0/0
Flow collector interface: cp-6/0/0
Interface state: Collecting flows
Destination change successful
```

Release Information

Command introduced before Junos OS Release 7.4.

request services flow-collector change-destination secondary interface

IN THIS SECTION

- Syntax | 349
- Description | 349
- Options | 349
- Required Privilege Level | 350
- Output Fields | 350
- Sample Output | 350

- Release Information | 350

Syntax

```
request services flow-collector change-destination secondary interface cp-fpc/pic/
port
<clear-files>
<clear-logs>
<immediately | gracefully>
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Switch to the secondary File Transfer Protocol (FTP) server that is configured as a flow collector.

Options

none	Switch to the secondary FTP server.
cp-*fpc/pic/port*	Use the specified flow collector interface name (**cp-*fpc/pic/port***) for the secondary destination.
clear-files	(Optional) Request clearing of existing data files in the FTP wait queue when the switch takes place.
clear-logs	(Optional) Request clearing of existing logs when the switch takes place.
immediately | gracefully	(Optional) Specify whether you want the switch to take place immediately, or to affect only newly created files.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services flow-collector change-destination secondary interface

```
user@host> request services flow-collector change-destination secondary interface
cp-6/0/0
Flow collector interface: cp-6/0/0
Interface state: Collecting flows
Destination change successful
```

Release Information

Command introduced before Junos OS Release 7.4.

request services flow-collector test-file-transfer

IN THIS SECTION

- Syntax | 351
- Description | 351

- Options | 351
- Required Privilege Level | 352
- Output Fields | 352
- Sample Output | 352
- Release Information | 352

Syntax

```
request services flow-collector test-file-transfer filename  interface (all | cp-fpc/pic/port)
(channel-zero | channel-one) (primary | secondary)
```

Description

(M40e, M160, and M320 Series routers, PTX Series, and T Series routers only) Transfer a test file to the primary or secondary File Transfer Protocol (FTP) server that is configured as a flow collector. This command verifies that the output side of the flow collector interface is operating properly.

Options

filename	Name of the test file to transfer.
interface (all | cp-*fpc/pic/port*)	Transfer a test file of flows from all configured flow collector interfaces or from only the specified interface.
channel-zero | channel-one	Transfer a file from export channel 0 (unit 0) or channel 1 (unit 1) of the PIC.
primary | secondary	Transfer a file to the primary or secondary server configured as a flow collector.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services flow-collector test-file-transfer interface channel-one primary

```
user@host> request services flow-collector test-file-transfer test_file interface cp-7/1/0
channel-one primary


Flow collector interface: cp-7/1/0
Interface state: Collecting flows
Response: Test file transfer successfully scheduled
```

Release Information

Command introduced before Junos OS Release 7.4.

show chassis forwarding

IN THIS SECTION

Syntax | 353

- Description | 353
- Options | 353
- Required Privilege Level | 353
- Output Fields | 354
- Sample Output | 354
- Release Information | 355

Syntax

```
show chassis forwarding
```

Description

Display status of the forwarding process (fwdd). This command is supported on Branch SRX Series Services Gateways.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 4 on page 354 lists the output fields for the `show chassis forwarding` command. Output fields are listed in the approximate order in which they appear.

Table 4: show chassis forwarding Output Fields

Field Name	Field Description

FWWD status	Forwarding status:
	• **State**:
	  • **Online**—FWDD is operational and running.
	  • **Offline**—FWDD is not running.
	• **Microkernel CPU utilization**—Percentage of microkernel CPU being used by the forwarding process.
	• **Real-time threads CPU utilization**—Percentage of CPU being used by the forwarding process.
	• **Heap utilization**—Percentage of heap space (dynamic memory) being used by the forwarding process. If this number exceeds 80 percent, there may be a software problem (memory leak).
	• **Buffer utilization**—Percentage of buffer space being used by the forwarding process for buffering internal messages.
	• **Uptime**—How long the forwarding process has been up and running.

Sample Output

show chassis forwarding

```
user@host> show chassis forwarding
FWDD status:
  State                                Online
  Microkernel CPU utilization        10 percent
  Real-time threads CPU utilization   4 percent
```

```
Heap utilization                    26 percent
Buffer utilization                   0 percent
Uptime:                               1 day, 1 hour, 30 minutes, 11 seconds
```

Release Information

Current—Command introduced before Junos OS Release 7.4.

Now—Command introduced in Junos OS Release 7.4. Support for Branch SRX Series added in Junos OS Release 10.1

show forwarding-options hyper-mode

IN THIS SECTION

- Syntax | 355
- Description | 356
- Required Privilege Level | 356
- Output Fields | 356
- Sample Output | 357
- Release Information | 357

Syntax

```
show forwarding-options hyper-mode
```

Description

Display information about the hyper mode feature. After you configure the hyper mode feature, you must reboot the system for the network device (a router or a switch) to reflect the change. For instance, if you have configured hyper mode but not rebooted the system, the `Current mode` field displays `normal` while the `Configured mode` field displays `hyper mode`.

> **NOTE**: Since hyper mode is a chassis-specific feature, the `show forwarding-options hyper-mode` command on a guest network function (GNF) displays the hyper-mode status of the chassis (BSYS), not of the GNF. To know more about GNFs, see Components of Junos Node Slicing.

Required Privilege Level

view

Output Fields

Table 5 on page 356 lists the output fields for the `show forwarding-options hyper-mode` command. Output fields are listed in the order in which they appear.

Table 5: show forwarding-options hyper-mode Output Fields

Field Name	Field Description
`Current mode`	Displays the current mode, either `normal` or `hyper mode`.
`Configured mode`	Displays the configured mode, either `normal` or `hyper mode`.

Sample Output

show forwarding-options hyper-mode

```
user@host> show forwarding-options hyper-mode
Current mode: hyper mode
Configured mode: hyper mode
```

Release Information

Command introduced in Junos OS Release 13.3R4.

RELATED DOCUMENTATION

- Configuring Hyper Mode on Enhanced MPCs to Speed Up Packet Processing | 124
- Understanding the Hyper Mode Feature on Enhanced MPCs for MX Series Routers and EX9200 Switches | 121
- hyper-mode (forwarding-options) | 219

show forwarding-options load-balance

IN THIS SECTION

- Syntax | 358
- Description | 358
- Options | 359
- Required Privilege Level | 359
- Output Fields | 359
- Sample Output | 360
- Sample Output | 361

Sample Output | 361
Sample Output | 361
Release Information | 362

Syntax

```
show forwarding-options load-balance
```

Description

Displays the load-balancing hash result. You can view this information for one, two, or three levels of load balancing. This command can be used in two ways to get different load-balancing information based on the parameters:

- To get the load-balancing decision result for routed IPv4, IPv6, and other L3 traffic, use the following:

  ```
  show forwarding-options load-balance ingress-interface <interface-name> family <family-type> source-address <src-IP> destination-address <dest-IP> transport-protocol <transport-protocol> source-port <src-port> destination-port <dest-port> tos <TOS>
  ```

  You can use this command when all the packet header details are available.

- To get the load-balancing decision result for raw packet dump files, use the following:

  ```
  show forwarding-options load-balance ingress-interface <interface-name> family <family-type> packet-dump <pkt-dump>
  ```

  You can use this command when the packets in the traffic are complex to be described by the 5-tuple.

NOTE:

- This feature is not supported for multicast flows and L2 packets.
- The maximum size of the raw packets entered on the CLI is 256 KB.

- When injecting raw packets, ensure 24 bytes of trailing space, which is accounted for in the packet header length. This space will be used for inserting metadata to the injected probe packet.

Options

interface-name	Ingress logical interface.
family-type	Layer 3 family “inet/inet6”.
src-IP	Source IP address.
dest-IP	Destination IP address.
transport-protocol	Transport layer protocol “tcp/udp”.
src-port	Source port (0 – 65535).
dest-port	Destination port (0 – 65535).
TOS	Type of service field (0 – 255).
pkt-dump	Raw packet dump in hexadecimal without '0x'. The hexadecimal code of packet information must start from the L3 header (IPv4/IPv6) to the end of the packet. It should be less than 256,000 characters in length.

Required Privilege Level

view

Output Fields

Table 6 on page 360 lists the output fields for the `show forwarding-options load-balance` command. Output fields are listed in the order in which they appear.

Table 6: show forwarding-options hyper-mode Output Fields

Field Name	Field Description
`Outgoing logical aggregate interface`	Egress aggregated Ethernet interface for current parameters
`Outgoing member physical interface`	Egress physical interface for current parameters
`Outgoing next hop address`	Egress next hop chosen for current parameters
`Outgoing next hop id`	Egress next-hop ID chosen for current parameters

Sample Output

AE Egress-IPv6 Egress

```
user@host> show forwarding-options load-balance family inet6 ingress-interface xe-5/0/3
transport-protocol tcp source-address 2201::2 destination-address 2202::2 source-port 1617
destination-port 1640 tos 224
================ fpc5 ================

Outgoing logical aggregate interface: ae0.0
Outgoing member physical interface  : xe-4/2/1
Outgoing next hop address           : fe80::2e21:72ff:fe71:1d
Outgoing next hop id                : 700
```

Sample Output

MPLS Egress IPv4

```
user@host> show forwarding-options load-balance ingress-interface xe-5/0/3 family inet source-
address 102.1.1.2 destination-address 202.1.1.2 tos 4 transport-protocol udp source-port 1431
destination-port 6531
================ fpc5 ================
Outgoing logical interface          : xe-5/0/2.0
Outgoing next hop address           : 40.40.40.2
Outgoing next hop id                : 747
```

Sample Output

Packet hash:IPv4 (MPLS egress)

```
user@host> show forwarding-options load-balance ingress-interface xe-5/0/3 family inet packet-
dump
4580004e000000003d114b1966010102ca0101020597198300 3a31abb9b7c07aa5bafa4349786960000000001101112134
bdfa97e00161a1b1c1d1e1f2021222324252627282 92a2b2c2d2e2fa073
================ fpc5 ================
Outgoing logical interface          : xe-4/1/1.0
Outgoing next hop address           : -
Outgoing next hop id                : 708
```

Note that for the MPLS egress, the outgoing next-hop address will not be displayed.

Sample Output

Packet hash:IPv6 (MPLS egress)

```
user@host> show forwarding-options load-balance ingress-interface xe-5/0/3 family inet6 packet-
dump
6c0000000026113d2101000000000000000000000000000022201000000000000000000000000000206b6078000265a60b
9b7c07aa5bafa434978696000000002101112133 61b3e6300021a1bd63a
```

```
================ fpc5 ================

Outgoing logical interface          : xe-4/1/2.0
Outgoing next hop address           : -
Outgoing next hop id                : 705
```

Release Information

Command introduced in Junos OS Release 17.1R1.

RELATED DOCUMENTATION

- show forwarding-options analyzer
- show forwarding-options port-mirroring | 362
- show forwarding-options hyper-mode | 355

show forwarding-options port-mirroring

IN THIS SECTION

- Syntax | 363
- Description | 363
- Options | 363
- Required Privilege Level | 363
- Output Fields | 363
- Sample Output | 364
- Release Information | 365

Syntax

```
show forwarding-options port-mirroring
<terse | detail>
<instance-name>
```

Description

Display current state of port-mirroring instances.

Options

terse | detail	(Optional) Display the specified level of output.
instance-name	(Optional) Display a single port-mirroring instance.

Required Privilege Level

view

Output Fields

Table 7 on page 363 lists the output fields for the `show forwarding-options port-mirroring` command. Output fields are listed in the approximate order in which they appear.

Table 7: show forwarding-options port-mirroring Output Fields

Field Name	Field Description	Level of Output
Instance Name	Name of port-mirroring instance.	All levels

Table 7: show forwarding-options port-mirroring Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Instance Id	Instance identification number.	All levels
State	Instance state, either `up` or `down`.	All levels
Input parameters		
Rate	Rate (ratio of packets sampled).	`detail`
Run-length	Run length (number of consecutive packets sampled).	`detail`
Maximum-packet-length	Maximum packet length.	`detail`
Output parameters		
Family	Protocol family.	`detail`
State	Instance state, either `up` or `down`.	`detail`
Destination	Destination (next-hop group name).	`detail`
Next-hop	IP address of the next hop to the destination.	`detail`

Sample Output

show forwarding-options port-mirroring terse

```
user@host> show forwarding-options port-mirroring terse
Instance Name                Instance Id    State
```

```
&global_instance             1            up
inst1                     2            up
```

show forwarding-options port-mirroring detail

```
user@host> show forwarding-options port-mirroring detail
Instance Name: pm1
  Instance Id: 2
  Input parameters:
    Rate                  : 2
    Run-length            : 0
    Maximum-packet-length : 0
  Output parameters:
    Family              State     Destination          Next-hop
    inet                up        ge-0/0/0.0           10.1.1.2
    inet6               up        ge-0/0/0.0           2001:db8::2
    any                 up        ge-0/0/1.0           NA
```

Release Information

Command introduced in Junos OS Release 9.6.

show forwarding-options next-hop-group

IN THIS SECTION

- Syntax | 366
- Description | 366
- Options | 366
- Required Privilege Level | 366
- Output Fields | 366
- Sample Output | 368

- Release Information | 370

Syntax

```
show forwarding-options next-hop-group
<terse | brief | detail>
<group-name>
```

Description

Display current state of next-hop groups.

Options

terse | brief | detail	(Optional) Display the specified level of output.
group-name	(Optional) Display a single next-hop group.

Required Privilege Level

view

Output Fields

Table 8 on page 367 lists the output fields for the `show forwarding-options next-hop-group` command. Output fields are listed in the approximate order in which they appear.

Table 8: show forwarding-options next-hop-group Output Fields

Field Name	Field Description	Level of Output
Next-hop-group	Name of next-hop group.	All levels
Type	Next-hop group type, such as **inet**, **inet6** or **layer-2**.	All levels
State	Next-hop group state, either **up** or **down**.	All levels
Members Interfaces	Names of interfaces to which next-hop group members belong.	**brief detail**
Member Subgroup	Names of subgroups to which next-hop group members belong.	**brief detail**
Number of members configured	Number of next-hop group members configured.	**detail**
Number of members that are up	Number of next-hop group members that are up.	**detail**
Number of subgroups configured	Number of subgroups configured.	**detail**
Number of subgroups that are up	Number of subgroups that are up.	**detail**

Sample Output

show forwarding-options next-hop-group terse

```
user@host> show forwarding-options next-hop-group terse
Next-hop-group                 Type             State
  nhg                            inet             up
  nhg6                           inet6            up
  vpls_nhg_2                           layer-2            down
```

show forwarding-options next-hop-group brief

```
user@host> show forwarding-options next-hop-group brief

Next-hop-group: nhg
      Type:  inet
      State: up
      Members Interfaces:
        ge-0/2/8.0             next-hop  192.0.2.10
        ge-5/1/8.0             next-hop  198.51.100.10
        ge-5/1/9.0             next-hop  203.0.113.10

Next-hop-group: nhg6
      Type:  inet6
      State: up
      Members Interfaces:
        ge-5/1/5.0             next-hop  2001:db8::1:10
        ge-5/1/6.0             next-hop  2001:db8::20:10        Member Subgroup: nhsg6
          Members Interfaces:
            ge-5/0/4.0             next-hop  2001:db8::3:1
            ge-5/1/4.0             next-hop  2001:db8::4:1

Next-hop-group: vpls_nhg_2
Type: layer-2        State: down
```

show forwarding-options next-hop-group detail

```
user@host> show forwarding-options next-hop-group detail

Next-hop-group: nhg
Type: inet
State: up
Number of members configured    : 3
Number of members that are up   : 3
Number of subgroups configured  : 0
Number of subgroups that are up : 0
Members Interfaces:                                 State
  ge-0/2/8.0             next-hop  192.0.2.10        up
  ge-5/1/8.0             next-hop  203.0.113.10        up
  ge-5/1/9.0             next-hop  198.51.100.10.10        up

Next-hop-group: nhg6
Type: inet6
State: up
Number of members configured    : 2
Number of members that are up   : 2
Number of subgroups configured  : 1
Number of subgroups that are up : 1
Members Interfaces:                                 State
ge-5/1/5.0             next-hop  2001:db8::1:10       up
ge-5/1/6.0             next-hop  2001:db8::20:10      up
Member Subgroup: nhsg6                            up
       Number of members configured    : 2
       Number of members that are up   : 2
       Members Interfaces:                                  State
           ge-5/0/4.0             next-hop  2001:db8::3:1   up
           ge-5/1/4.0             next-hop  2001:db8::4:1   up

Next-hop-group: vpls_nhg_2
Number of members configured    : 2
Number of members that are up   : 0
Number of subgroups configured  : 0
Number of subgroups that are up : 0
Type: layer-2        State: down
Members Interfaces:             State
```

```
ge-2/2/1.100                down
ge-2/3/9.0                  down
```

Release Information

Command introduced in Junos OS Release 9.6.

Support for IPv6 introduced in Junos OS Release 14.2 for the MX Series routers.

RELATED DOCUMENTATION

show forwarding-options port-mirroring

show interfaces (Flow Monitoring)

IN THIS SECTION

- Syntax | **370**
- Description | **371**
- Options | **371**
- Required Privilege Level | **371**
- Output Fields | **371**
- Sample Output | **377**
- Release Information | **378**

Syntax

```
show interfaces mo-fpc/pic/port:channel
<brief | detail | extensive | terse>
<descriptions>
```

```
<media>
<snmp-index snmp-index>
<statistics>
```

Description

(M Series and T Series routers only) Display status information about the specified flow monitoring interface.

Options

mo-*fpc/pic/port*:*channel*	Display standard status information about the specified flow monitoring interface.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 9 on page 372 lists the output fields for the `show interfaces` (Flow Monitoring) command. Output fields are listed in the approximate order in which they appear.

Table 9: show interfaces Output Fields (Flow Monitoring)

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Link	Status of the link: **up** or **down**.	All levels
Enabled	State of the interface. Possible values are described in the "Enabled Field" section under *Common Output Fields Description*.	All levels
Interface index	Physical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Description	Description and name of the interface.	All levels
Type	Type of interface.	All levels
Link-level type	Encapsulation type used on the physical interface.	All levels
MTU	Maximum Transmit Unit (MTU). Size of the largest packet to be transmitted.	All levels
Clocking	Reference clock source of the interface.	All levels
Speed	Network speed on the interface.	All levels

Table 9: show interfaces Output Fields (Flow Monitoring) *(Continued)*

Field Name	Field Description	Level of Output
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under *Common Output Fields Description*.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under *Common Output Fields Description*.	All levels
Link type	Data transmission type.	All levels
Link flags	Information about the link. Possible values are described in the "Link Flags" section under *Common Output Fields Description*.	All levels
Physical info	Information about the physical interface.	All levels
Hold-times	Current interface hold-time up and hold-time down. Value is in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	Media access control (MAC) address of the interface.	**detail extensive** none
Alternate link address	Backup link address.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped: *year-month-day hour.minute.second timezone* (*hour.minute.second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**	**detail extensive**
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**

Table 9: show interfaces Output Fields (Flow Monitoring) *(Continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**, **Output bytes**—Number of bytes received and transmitted on the interface. • **Input packets**, **Output packets**—Number of packets received and transmitted on the interface.	**detail extensive**
Input errors	• **Errors**—Input errors on the interface. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Frames received smaller than the runt threshold. • **Giants**—Frames received larger than the giant threshold. • **Policed Discards**—Frames that the incoming packet match code discarded because the frames did not recognize them or were not of interest. Usually, this field reports protocols that Junos does not support. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 9: show interfaces Output Fields (Flow Monitoring) *(Continued)*

Field Name	Field Description	Level of Output
Output errors	• **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly, possibly once every 10 seconds, the cable, the remote system, or the interface is malfunctioning. • **Errors**—Sum of outgoing frame terminates and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet dropped by the ASIC Red mechanism. • **Resource errors**—Sum of transmit drops.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface; values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels

Table 9: show interfaces Output Fields (Flow Monitoring) *(Continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface. • **Input packets, Output packets**—Number of packets received and transmitted on the interface.	**detail extensive**
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	**detail extensive**
Protocol	Protocol family configured on the logical interface (such as **iso** or **inet6**).	**detail extensive** none
MTU	MTU size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Route table in which this address exists; for example, **Route table:0** refers to inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under *Common Output Fields Description*.	**detail extensive** none

Sample Output

show interfaces extensive (Flow Monitoring)

```
user@host> show interfaces mo-4/0/0     extensive
Physical interface: mo-4/0/0, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 42, Generation: 28
  Description: monitor pic 2
  Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: Unlimited,
  Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : 2005-05-24 16:43:12 PDT (00:17:46 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :            756824218              8328536 bps
   Output bytes  :            872916185              8400160 bps
   Input  packets:               508452                  697 pps
   Output packets:             15577196                18750 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 2, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

  Logical interface mo-4/0/0.0 (Index 83) (SNMP ifIndex 43) (Generation 26)
    Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
    Traffic statistics:
     Input  bytes  :            756781796
     Output bytes  :            872255328
     Input  packets:               507233
     Output packets:             15575988
    Local statistics:
     Input  bytes  :                   0
```

```
     Output bytes  :                    0
     Input  packets:                    0
     Output packets:                    0
    Transit statistics:
     Input  bytes  :            756781796              8328536 bps
     Output bytes  :            872255328              8400160 bps
     Input  packets:               507233                  697 pps
     Output packets:             15575988                18750 pps
    Protocol inet, MTU: Unlimited, Generation: 38, Route table: 0
      Flags: None

  Logical interface mo-4/0/0.16383 (Index 84) (SNMP ifIndex 58) (Generation 27)
...
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces (M Series, MX Series, T Series Routers, and PTX Series Management and Internal Ethernet)

IN THIS SECTION

- Syntax (M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface) | **379**
- Syntax (M Series, MX Series, T Series, and PTX Series Routers Internal Ethernet Interface) | **379**
- Description | **379**
- Options | **380**
- Required Privilege Level | **380**
- Output Fields | **380**
- Sample Output | **386**

- Release Information | 403

Syntax (M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface)

```
show interfaces em0 | fxp0 | mgmtre0
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Syntax (M Series, MX Series, T Series, and PTX Series Routers Internal Ethernet Interface)

```
show interfaces bcm0  | em0 | em1| fxp1 | fxp2 | ixgbe0 | ixgbe1
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

(M Series, T Series, TX Matrix Plus, and PTX Series devices only) Display status information about the management Ethernet and internal Ethernet interfaces.

Options

em0 | fxp0 | mgmtre0	(M Series, MX Series, T Series, and PTX Series) Display standard information about the management Ethernet interface. For supported Ethernet interface by chassis and Routing Engine, see Supported Routing Engines by Router.
bcm0 | em0 | em1 | fxp1 | fxp2 | ixgbe0 | ixgbe1	(M Series, MX Series, T Series, and PTX Series) Display standard information about the internal Ethernet interfaces. See Supported Routing Engines by Router for the internal Ethernet interface names for each Routing Engine by hardware platform. **NOTE**: On Junos OS Evolved, the ixgbe0 and ixgbe1 internal interfaces are deprecated.
brief | detail | extensive | terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information.
snmp-index *snmp-index*	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

Table 10 on page 381 lists the output fields for the `show interfaces` (management) command on the M Series routers, T Series routers, TX Matrix Plus routers, and PTX Series. Output fields are listed in the approximate order in which they appear.

Table 10: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the "Enabled Field" section under *Common Output Fields Description*.	All levels
Interface index	Physical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Type	Type of interface.	All levels
Link-level type	Encapsulation type used on the physical interface.	All levels
MTU	Maximum transmission unit (MTU)—Size of the largest packet to be transmitted.	All levels
Clocking	Reference clock source of the interface.	All levels
Speed	Network speed on the interface.	All levels
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under *Common Output Fields Description*.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under *Common Output Fields Description*.	All levels

Table 10: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface *(Continued)*

Field Name	Field Description	Level of Output
Link type	Data transmission type.	**detail extensive** none
Link flags	Information about the link. Possible values are described in the "Link Flags" section under *Common Output Fields Description*.	**detail extensive**
Physical info	Information about the physical interface.	**detail extensive**
Hold-times	Current interface hold-time up and hold-time down. Value is in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	Media access control (MAC) address of the interface.	**detail extensive** none
Alternate link address	Backup link address.	**detail extensive**
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped: *year-month-day hour:minute:second timezone* (*hour:minute:second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.	**detail extensive** none
Input packets	Number of packets received on the physical interface.	None specified
Output packets	Number of packets transmitted on the physical interface.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**

Table 10: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface *(Continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the logical and physical interface. • **Input bytes**, **Output bytes**—Number of bytes received and transmitted on the interface. • **Input packets**, **Output packets**—Number of packets received and transmitted on the interface.	**detail extensive**
Input errors	• **Errors**—Input errors on the interface. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Frames received smaller than the runt threshold. • **Giants**—Frames received larger than the giant threshold. • **Policed Discards**—Frames that the incoming packet match code discarded because they were not recognized or were not of interest. Usually, this field reports protocols that Junos does not support. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 10: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface *(Continued)*

Field Name	Field Description	Level of Output
Output errors	• **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly, possibly once every 10 seconds, the cable, the remote system, or the interface is malfunctioning. • **Errors**—Sum of outgoing frame terminates and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet dropped by the ASIC RED mechanism. • **Resource errors**—Sum of transmit drops.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface	All levels
Index	Logical interface index number, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	Logical interface SNMP interface index number.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface; values are described in the "Device Flags" section under *Common Output Fields Description*.	All levels
Encapsulation	Encapsulation on the logical interface.	**detail extensive** none
inet	IP address of the logical interface.	**brief**

Table 10: show interfaces Output Fields for M Series, MX Series, T Series, and PTX Series Routers Management Ethernet Interface *(Continued)*

Field Name	Field Description	Level of Output
Protocol	Protocol family configured on the logical interface (such as **iso** or **inet6**).	**detail extensive** none
MTU	MTU size on the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route table	Route table in which this address exists. For example, **Route table:0** refers to inet.0.	**detail extensive**
Flags	Information about the protocol family flags. Possible values are described in the "Family Flags" section under *Common Output Fields Description*.	**detail extensive** none
Addresses, Flags	Information about address flags. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces brief (Management Ethernet)

```
user@host> show interfaces fxp0 brief
Physical interface: fxp0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface fxp0.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  192.168.70.143/21
```

show interfaces (Management Ethernet)

```
user@host> show interfaces fxp0
Physical interface: fxp0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 1
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Half-Duplex
  Current address: 00:00:5E:00:53:89, Hardware address: 00:00:5E:00:53:89
  Last flapped   : Never
    Input packets : 80804
    Output packets: 1105

  Logical interface fxp0.0 (Index 2) (SNMP ifIndex 13)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.64/21, Local: 192.168.70.143,
        Broadcast: 192.168.71.255
```

show interfaces (Management Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces em0
Physical interface: em0, Enabled, Physical link is Up
  Interface index: 8, SNMP ifIndex: 17
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:00:5E:00:53:c0, Hardware address: 00:00:5E:00:53:c0
  Last flapped   : Never
    Input packets : 1424
    Output packets: 5282

  Logical interface em0.0 (Index 3) (SNMP ifIndex 18)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 1424
    Output packets: 5282
    Protocol inet, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.178.0/25, Local: 192.168.178.11, Broadcast: 192.168.178.127
```

show interfaces (Management Ethernet [PTX Series Packet Transport Routers])

```
user@host> show interfaces em0
Physical interface: em0, Enabled, Physical link is Up
      Interface index: 8, SNMP ifIndex: 0
      Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
      Device flags   : Present Running
      Interface flags: SNMP-Traps
      Link type      : Full-Duplex
      Current address: 00:00:5E:00:53:1b, Hardware address: 00:00:5E:00:53:1b
      Last flapped   : Never
    Input packets : 212581
    Output packets: 71

      Logical interface em0.0 (Index 3) (SNMP ifIndex 0)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 212551
```

```
    Output packets: 71
    Protocol inet, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 192.168.3/24, Local: 192.168.3.30,
        Broadcast: 192.168.3.255
```

show interfaces detail (Management Ethernet)

```
user@host> show interfaces fxp0 detail
Physical interface: fxp0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 1, Generation: 0
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Half-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:89, Hardware address: 00:00:5E:00:53:89
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              6484031
   Output bytes  :               167503
   Input  packets:                81008
   Output packets:                 1110

  Logical interface fxp0.0 (Index 2) (SNMP ifIndex 13) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 6, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.64/21, Local: 192.168.70.143,
        Broadcast: 192.168.71.255, Generation: 1
```

show interfaces detail (Management Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces em0 detail
Physical interface: em0, Enabled, Physical link is Up
  Interface index: 8, SNMP ifIndex: 17, Generation: 2
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:c0, Hardware address: 00:00:5E:00:53:c0
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :               124351
   Output bytes  :              1353212
   Input  packets:                 1804
   Output packets:                 5344
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0

  Logical interface em0.0 (Index 3) (SNMP ifIndex 18) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :               117135
     Output bytes  :              1331647
     Input  packets:                 1804
     Output packets:                 5344
    Local statistics:
     Input  bytes  :               117135
     Output bytes  :              1331647
     Input  packets:                 1804
     Output packets:                 5344
    Protocol inet, MTU: 1500, Generation: 1, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
```

```
        Destination: 192.168.178.0/25, Local: 192.168.178.11, Broadcast: 192.168.178.127,
Generation: 1
```

show interfaces detail (Management Ethernet [PTX Packet Transport Routers])

```
user@host> show interfaces detail em0
Physical interface: em0, Enabled, Physical link is Up
      Interface index: 8, SNMP ifIndex: 0, Generation: 3
      Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
      Speed: 1000mbps
      Device flags   : Present Running
      Interface flags: SNMP-Traps
      Link type      : Full-Duplex
      Physical info  : Unspecified
      Hold-times     : Up 0 ms, Down 0 ms
      Current address: 00:00:5E:00:53:1b, Hardware address: 00:00:5E:00:53:1b
      Alternate link address: Unspecified
      Last flapped   : Never
      Statistics last cleared: Never
      Traffic statistics:
       Input  bytes  :             15255909
       Output bytes  :                 4608
       Input  packets:               214753
       Output packets:                   72
       IPv6 transit statistics:
    Input  bytes  :                    0
    Output bytes  :                    0
    Input  packets:                    0
    Output packets:                    0

      Logical interface em0.0 (Index 3) (SNMP ifIndex 0) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :             14394630
     Output bytes  :                 3024
     Input  packets:               214723
     Output packets:                   72
    Local statistics:
     Input  bytes  :             14394630
     Output bytes  :                 3024
     Input  packets:               214723
```

```
     Output packets:                  72
    Protocol inet, MTU: 1500, Generation: 1, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 192.168.3/24, Local: 192.168.3.30,
        Broadcast: 192.168.3.255, Generation: 1
```

show interfaces extensive (Management Ethernet)

```
user@host> show interfaces fxp0 extensive
Physical interface: fxp0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 1, Generation: 0
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Half-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:89, Hardware address: 00:00:5E:00:53:89
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              6678904
   Output bytes  :               169657
   Input  packets:                83946
   Output packets:                 1127
  Input errors:
    Errors: 12, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

  Logical interface fxp0.0 (Index 2) (SNMP ifIndex 13) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 6, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
```

```
        Destination: 192.168.64/21, Local: 192.168.70.143,
        Broadcast: 192.168.71.255, Generation: 1
```

show interfaces extensive (Management Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces em0 extensive

Physical interface: em0, Enabled, Physical link is Up
  Interface index: 8, SNMP ifIndex: 17, Generation: 2
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:c0, Hardware address: 00:00:5E:00:53:c0
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :               127120
   Output bytes  :              1357414
   Input  packets:                 1843
   Output packets:                 5372
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0

  Logical interface em0.0 (Index 3) (SNMP ifIndex 18) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :               119748
     Output bytes  :              1335719
     Input  packets:                 1843
```

```
     Output packets:                 5372
    Local statistics:
     Input  bytes  :               119748
     Output bytes  :              1335719
     Input  packets:                 1843
     Output packets:                 5372
    Protocol inet, MTU: 1500, Generation: 1, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.168.178.0/25, Local: 192.168.178.11, Broadcast: 192.168.178.127,
Generation: 1
```

show interfaces extensive (Management Ethernet [PTX Series Packet Transport Routers])

```
user@host> show interfaces extensive em0
Physical interface: em0, Enabled, Physical link is Up
      Interface index: 8, SNMP ifIndex: 0, Generation: 3
      Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
      Speed: 1000mbps
      Device flags   : Present Running
      Interface flags: SNMP-Traps
      Link type      : Full-Duplex
      Physical info  : Unspecified
      Hold-times     : Up 0 ms, Down 0 ms
      Current address: 00:00:5E:00:53:1b, Hardware address: 00:00:5E:00:53:1b
      Alternate link address: Unspecified
      Last flapped   : Never
      Statistics last cleared: Never
      Traffic statistics:
       Input  bytes  :             15236459
       Output bytes  :                 4608
       Input  packets:               214482
       Output packets:                   72
       IPv6 transit statistics:
    Input  bytes  :                    0
    Output bytes  :                    0
    Input  packets:                    0
    Output packets:                    0
      Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
```

```
      Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

      Logical interface em0.0 (Index 3) (SNMP ifIndex 0) (Generation 1)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :             14376264
     Output bytes  :                 3024
     Input  packets:               214452
     Output packets:                   72
    Local statistics:
     Input  bytes  :             14376264
     Output bytes  :                 3024
     Input  packets:               214452
     Output packets:                   72
    Protocol inet, MTU: 1500, Generation: 1, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 192.168.3/24, Local: 192.168.3.30,
        Broadcast: 192.168.3.255, Generation: 1
```

show interfaces mgmtre0 (Management Ethernet [PTX5000 Router])

```
user@host> show interfaces mgmtre0 extensive
Physical interface: mgmtre0, Enabled, Physical link is Up
  Interface index: 1001, SNMP ifIndex: 501
  Link-level type: Ethernet, MTU: 1500
  Device flags   : Present
  Interface flags: None
  Link flags     : None
  Current address: ec:9e:cd:06:30:da, Hardware address: ec:9e:cd:06:30:da
  Last flapped   : Never

  Logical interface mgmtre0.0 (Index 1001) (SNMP ifIndex 503)
    Flags: Encapsulation: ENET2
    Protocol inet, MTU: 1486
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.92.248/23, Local: 10.92.248.22,
        Broadcast: 10.92.249.255
```

```
    Protocol multiservice, MTU: Unlimited
      Flags: None
```

show interfaces brief (Management Ethernet)

```
user@host> show interfaces fxp1 brief
Physical interface: fxp1, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface fxp1.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  10.0.0.4/8
    inet6 fe80::200:ff:fe00:4/64
          fec0::10:0:0:4/64
    tnp   4
```

show interfaces brief (Management Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces em0 brief
Physical interface: em0, Enabled, Physical link is Up
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps

  Logical interface em0.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  192.168.178.11/25
```

show interfaces brief (Management Ethernet [PTX Series Packet Transport Routers])

```
user@host> show interfaces em0 brief
Physical interface: em0, Enabled, Physical link is Up
      Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
      Speed: 1000mbps
      Device flags   : Present Running
```

```
      Interface flags: SNMP-Traps

      Logical interface em0.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  192.168.3.30/24

    root@aboslutely> show interfaces em0 terse
    Interface               Admin Link Proto    Local                 Remote
    em0                     up    up
    em0.0                   up    up   inet     192.168.3.30/24
```

show interfaces (Internal Ethernet)

```
user@host> show interfaces fxp1
Physical interface: fxp1, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 2
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Last flapped   : Never
    Input packets : 30655
    Output packets: 33323

  Logical interface fxp1.0 (Index 3) (SNMP ifIndex 14)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10/8, Local: 10.0.0.4, Broadcast: 10.255.255.255
    Protocol inet6, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe00:4
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::10:0:0:4
    Protocol tnp, MTU: 1500
      Flags: Primary, Is-Primary
```

```
      Addresses
        Local: 4
```

show interfaces (Internal Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces ixgbe0
Physical interface: ixgbe0, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 116
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Last flapped   : Never
    Input packets : 2301738
    Output packets: 3951155

  Logical interface ixgbe0.0 (Index 4) (SNMP ifIndex 117)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 2301595
    Output packets: 3951155
    Protocol inet, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: 10/8, Local: 10.34.0.4, Broadcast: 10.255.255.255
      Addresses, Flags: Primary Is-Default Is-Preferred Is-Primary
        Destination: 192.168/16, Local: 192.168.0.4, Broadcast: 192.168.0.4
    Protocol inet6, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe22:4
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::a:22:0:4
    Protocol tnp, MTU: 1500
      Flags: Primary, Is-Primary
      Addresses
        Local: 0x22000004
```

show interfaces detail (Internal Ethernet)

```
user@host> show interfaces fxp1 detail
Physical interface: fxp1, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 2, Generation: 1
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              2339969
   Output bytes  :             15880707
   Input  packets:                30758
   Output packets:                33443

  Logical interface fxp1.0 (Index 3) (SNMP ifIndex 14) (Generation 2)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 7, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10/8, Local: 10.0.0.4, Broadcast: 10.255.255.255,
        Generation: 3
    Protocol inet6, MTU: 1500, Generation: 8, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe00:4,
        Broadcast: Unspecified, Generation: 5
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::10:0:0:4, Broadcast: Unspecified,
        Generation: 7
    Protocol tnp, MTU: 1500, Generation: 9, Route table: 1
      Flags: Primary, Is-Primary
      Addresses, Flags: None
```

```
        Destination: Unspecified, Local: 4, Broadcast: Unspecified,
        Generation: 8
```

show interfaces detail (Internal Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces ixgbe0 detail
Physical interface: ixgbe0, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 116, Generation: 3
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :            238172825
   Output bytes  :           1338948955
   Input  packets:              2360984
   Output packets:              4061512
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0

  Logical interface ixgbe0.0 (Index 4) (SNMP ifIndex 117) (Generation 2)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :            228720309
     Output bytes  :           1261387447
     Input  packets:              2360841
     Output packets:              4061512
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
```

```
    Local statistics:
     Input  bytes  :            228720309
     Output bytes  :           1261387447
     Input  packets:              2360841
     Output packets:              4061512
    Protocol inet, MTU: 1500, Generation: 2, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: 10/8, Local: 10.34.0.4, Broadcast: 10.255.255.255, Generation: 2
      Addresses, Flags: Primary Is-Default Is-Preferred Is-Primary
        Destination: 192.168/16, Local: 192.168.0.4, Broadcast: 191.255.255.255, Generation: 3
    Protocol inet6, MTU: 1500, Generation: 3, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe22:4
    Generation: 4
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::a:22:0:4
    Protocol tnp, MTU: 1500, Generation: 5
    Generation: 4, Route table: 1
      Flags: Primary, Is-Primary
      Addresses, Flags: None
        Destination: Unspecified, Local: 0x22000004, Broadcast: Unspecified, Generation: 6
```

show interfaces extensive (internal Ethernet)

```
user@host> show interfaces fxp1 extensive
Physical interface: fxp1, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 2, Generation: 1
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified,
  Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
```

```
   Input  bytes  :              2349897
   Output bytes  :             15888605
   Input  packets:                30896
   Output packets:                33607
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

  Logical interface fxp1.0 (Index 3) (SNMP ifIndex 14) (Generation 2)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 7, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10/8, Local: 10.0.0.4, Broadcast: 10.255.255.255,
        Generation: 3
    Protocol inet6, MTU: 1500, Generation: 8, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::200:ff:fe00:4,
        Broadcast: Unspecified, Generation: 5
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: fec0::/64, Local: fec0::10:0:0:4, Broadcast: Unspecified,
        Generation: 7
    Protocol tnp, MTU: 1500, Generation: 9, Route table: 1
      Flags: Primary, Is-Primary
      Addresses, Flags: None
        Destination: Unspecified, Local: 4, Broadcast: Unspecified,
        Generation: 8
```

show interfaces extensive (internal Ethernet [TX Matrix Plus Router])

```
user@host> show interfaces ixgbe0 extensive
Physical interface: ixgbe0, Enabled, Physical link is Up
  Interface index: 2, SNMP ifIndex: 116, Generation: 3
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Clocking: Unspecified, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
```

```
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:04, Hardware address: 00:00:5E:00:53:04
  Alternate link address: Unspecified
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :            242730780
   Output bytes  :           1348312269
   Input  packets:              2398737
   Output packets:              4133510
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0

  Logical interface ixgbe0.0 (Index 4) (SNMP ifIndex 117) (Generation 2)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :            233127252
     Output bytes  :           1269350897
     Input  packets:              2398594
     Output packets:              4133510
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Local statistics:
     Input  bytes  :            233127252
     Output bytes  :           1269350897
     Input  packets:              2398594
     Output packets:              4133510
    Protocol inet, MTU: 1500, Generation: 2, Route table: 1
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: 10/8, Local: 10.34.0.4, Broadcast: 10.255.255.255, Generation: 2
```

```
    Addresses, Flags: Primary Is-Default Is-Preferred Is-Primary
      Destination: 192.168/16, Local: 192.168.0.4, Broadcast: 191.255.255.255, Generation: 3
  Protocol inet6, MTU: 1500, Generation: 3, Route table: 1
    Flags: Is-Primary
    Addresses, Flags: Is-Preferred
      Destination: fe80::/64, Local: fe80::200:ff:fe22:4
  Generation: 4
    Addresses, Flags: Is-Default Is-Preferred Is-Primary
      Destination: fec0::/64, Local: fec0::a:22:0:4
  Protocol tnp, MTU: 1500, Generation: 5
  Generation: 4, Route table: 1
    Flags: Primary, Is-Primary
    Addresses, Flags: None
      Destination: Unspecified, Local: 0x22000004, Broadcast: Unspecified, Generation: 6
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced on PTX Series Packet Transport Routers for Junos OS Release 12.1.

show interfaces statistics

IN THIS SECTION

- Syntax | **404**
- Description | **404**
- Options | **404**
- Required Privilege Level | **405**
- Output Fields | **405**
- Sample Output | **405**
- Release Information | **422**

Syntax

```
show interfaces statistics interface-name
<satellite-device [device-alias-name |all ]>
<detail>
```

Description

Display static interface statistics, such as errors.

NOTE: When the **show interfaces statistics** command is executed on an interface that is configured on T4000 Type 5 FPC, the *IPv6 transit statistics* field displays:

- Total statistics (sum of transit and local statistics) at the physical interface level
- Transit statistics at the logical interface level

Options

interface-name	Name of an interface.
satellite-device [*device-alias-name* | all]	(Junos Fusion only) (Optional) Display interface statistics for interfaces on the specified satellite device in the Junos Fusion, or on all satellite devices in the Junos Fusion.
	NOTE: In a Junos Fusion Enterprise, logical interface statistics are not synced across aggregation devices in a dual-aggregation device topology.
detail	(Optional) Display detailed output.

Required Privilege Level

view

Output Fields

Output from both the `show interfaces` *`interface-name`* `detail` and the `show interfaces` *`interface-name`* `extensive` commands include all the information displayed in the output from the `show interfaces statistics` command. For more information, see the particular interface type in which you are interested. For information about destination class and source class statistics, see the “Destination Class Field” section and the “Source Class Field” section under Common Output Fields Description. For information about the input errors and output errors, see Fast Ethernet and Gigabit Ethernet Counters.

Sample Output

show interfaces statistics (Fast Ethernet)

```
user@host> show interfaces fe-1/3/1 statistics
Physical interface: fe-1/3/1, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 1042
  Description: ford fe-1/3/1
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:00:5E:00:53:dc, Hardware address: 00:00:5E:00:53:dc
  Last flapped   : 2006-04-18 03:08:59 PDT (00:01:24 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0
  Active alarms  : None
  Active defects : None
  Logical interface fe-1/3/1.0 (Index 69) (SNMP ifIndex 50)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500
```

```
      Flags: Is-Primary, DCU, SCU-in
                                              Packets                Bytes
      Destination class            (packet-per-second)    (bits-per-second)
                          silver1                    0                    0
                                  (                  0) (                 0)
                          silver2                    0                    0
                                  (                  0) (                 0)
                          silver3                    0                    0
                                  (                  0) (                 0)
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 10.27.245/24, Local: 10.27.245.2,
        Broadcast: 10.27.245.255
    Protocol iso, MTU: 1497
      Flags: Is-Primary
```

show interfaces statistics (Gigabit Ethernet PIC—Egress)

```
user@host> show interfaces ge-5/2/0 statistics detail
Physical interface: ge-5/2/0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 519, Generation: 149
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, BPDU Error: None, MAC-REWRITE Error:
None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled, Remote fault:
Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5E:00:53:74, Hardware address: 00:00:5E:00:53:74
  Last flapped   : 2009-11-11 11:24:00 PST (09:23:08 ago)
  Statistics last cleared: 2009-11-11 17:50:58 PST (02:56:10 ago)
  Traffic statistics:
   Input  bytes  :               271524                   0 bps
   Output bytes  :             37769598                 352 bps
   Input  packets:                 3664                   0 pps
   Output packets:               885790                   0 pps
   IPv6 transit statistics:
    Input  bytes  :                    0
    Output bytes  :             16681118
    Input  packets:                    0
```

```
    Output packets:              362633
   Multicast statistics:
    IPV4 multicast statistics:
     Input  bytes  :               112048                   0 bps
     Output bytes  :             20779920                   0 bps
     Input  packets:                 1801                   0 pps
     Output packets:               519498                   0 pps
    IPV6 multicast statistics:
     Input  bytes  :               156500                   0 bps
     Output bytes  :             16681118                   0 bps
     Input  packets:                 1818                   0 pps
     Output packets:               362633                   0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0,
    L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0,
    Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort               882558               882558                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                3232                 3232                    0
  Active alarms  : None
  Active defects : None

  Logical interface ge-5/2/0.0 (Index 71) (SNMP ifIndex 573) (Generation 135)
    Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
    Egress account overhead: 100
    Ingress account overhead: 90
    Traffic statistics:
     Input  bytes  :               271524
     Output bytes  :             37769598
     Input  packets:                 3664
     Output packets:               885790
     IPv6 transit statistics:
      Input  bytes  :                    0
      Output bytes  :             16681118
      Input  packets:                    0
      Output packets:               362633
```

```
    Local statistics:
     Input  bytes  :               271524
     Output bytes  :               308560
     Input  packets:                 3664
     Output packets:                 3659
    Transit statistics:
     Input  bytes  :                    0                    0 bps
     Output bytes  :             37461038                    0 bps
     Input  packets:                    0                    0 pps
     Output packets:               882131                    0 pps
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :            16681118
      Input  packets:                   0
      Output packets:              362633
    Multicast statistics:
     IPV4 multicast statistics:
      Input  bytes  :               112048                    0 bps
      Output bytes  :             20779920                    0 bps
      Input  packets:                1801                    0 pps
      Output packets:              519498                    0 pps
     IPV6 multicast statistics:
      Input  bytes  :               156500                    0 bps
      Output bytes  :             16681118                    0 bps
      Input  packets:                1818                    0 pps
      Output packets:              362633                    0 pps
    Protocol inet, MTU: 1500, Generation: 151, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.40.40.0/30, Local: 10.40.40.2, Broadcast: 10.40.40.3, Generation: 167
    Protocol inet6, MTU: 1500, Generation: 152, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: ::10.40.40.0/126, Local: ::10.40.40.2
    Generation: 169
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::21d:b5ff:fe61:d974
    Protocol multiservice, MTU: Unlimited, Generation: 171
    Generation: 153, Route table: 0
      Policer: Input: __default_arp_policer__
```

show interfaces statistics detail (Aggregated Ethernet)

```
user@host> show interfaces ae0 detail
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 186, SNMP ifIndex: 111, Generation: 187
  Link-level type: Ethernet, MTU: 1514, Speed: 2000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1,
  Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:00:5E:0053:f0, Hardware address: 00:00:5E:00:53:f0
  Last flapped   : Never
  Statistics last cleared: 2006-12-23 03:04:16 PST (01:16:24 ago)
  Traffic statistics:
   Input  bytes  :                28544                    0 bps
   Output bytes  :                39770                    0 bps
   Input  packets:                  508                    0 pps
   Output packets:                  509                    0 pps
   Input  bytes  :           IPv6 28544
   Output bytes  :               IPv6 0
   Input  packets:             IPv6 508
   Output packets:               IPv6 0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

  Logical interface ae0.0 (Index 67) (SNMP ifIndex 139) (Generation 145)
    Flags: SNMP-Traps Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :           508          0         28544            0
        Output:           509          0         35698            0
    Link:
      ge-3/3/8.0
        Input :           508          0         28544            0
        Output:             0          0             0            0
      ge-3/3/9.0
        Input :             0          0             0            0
        Output:             0          0             0            0
```

```
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-3/3/8.0                0           0            0            0
      ge-3/3/9.0                0           0            0            0
  Egress queues: 8 supported, 8 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                   0                    0                    0
    Protocol inet, MTU: 1500, Generation: 166, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: 10.1.1.255,
        Generation: 159
    Protocol inet6, MTU: 1500, Generation: 163, Route table: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::206:5bff:fe05:c321,
        Broadcast: Unspecified, Generation: 161
```

show interfaces statistics detail (Aggregated Ethernet—Ingress)

```
user@host> show interfaces statistics detail ae0 | no-more
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 504, Generation: 278
  Link-level type: Ethernet, MTU: 1514, Speed: 1Gbps, BPDU Error: None, MAC-REWRITE Error: None,
Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth
needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:00:5E:00:53:f0, Hardware address: 00:00:5E:00:53:f0
  Last flapped   : 2009-11-09 03:30:23 PST (00:01:28 ago)
  Statistics last cleared: 2009-11-09 03:26:18 PST (00:05:33 ago)
  Traffic statistics:
   Input  bytes  :            544009602             54761856 bps
   Output bytes  :                 3396                    0 bps
   Input  packets:             11826292               148809 pps
   Output packets:                   42                    0 pps
   IPv6 transit statistics:
    Input  bytes  :           350818604
```

```
    Output bytes  :                   0
    Input  packets:             7626488
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Ingress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                   0                    0                    0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                   21                   21                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                 451                  451                    0

  Logical interface ae0.0 (Index 70) (SNMP ifIndex 574) (Generation 177)
    Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :      11826292     148809     544009602     54761856
        Output:            42          0          3396            0
    Link:
      ge-5/2/0.0
        Input :      11826292     148809     544009602     54761856
        Output:            42          0          3396            0
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-5/2/0.0                 0           0            0            0
    Protocol inet, MTU: 1500, Generation: 236, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.30.30.0/30, Local: 10.30.30.2, Broadcast: 10.30.30.3, Generation: 310
    Protocol inet6, MTU: 1500, Generation: 237, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: ::10.30.30.0/126, Local: ::10.30.30.2
    Generation: 312
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::21d:b5ff:fe61:dbf0
    Protocol multiservice, MTU: Unlimited, Generation: 314
```

```
    Generation: 238, Route table: 0
      Policer: Input: __default_arp_policer__
```

show interfaces statistics detail (Aggregated Ethernet—Egress)

```
user@host> show interfaces statistics detail ae0 | no-more
Physical interface: ae0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 501, Generation: 319
  Link-level type: Ethernet, MTU: 1514, Speed: 1Gbps, BPDU Error: None, MAC-REWRITE Error: None,
Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1, Minimum bandwidth
needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 00:00:5E:00:53:f0, Hardware address: 00:00:5E:00:53:f0
  Last flapped   : 2009-11-09 03:30:24 PST (00:02:42 ago)
  Statistics last cleared: 2009-11-09 03:26:42 PST (00:06:24 ago)
  Traffic statistics:
   Input  bytes  :                 440                   0 bps
   Output bytes  :          1047338120            54635848 bps
   Input  packets:                   7                   0 pps
   Output packets:            22768200              148466 pps
   IPv6 transit statistics:
    Input  bytes  :                 288
    Output bytes  :           723202616
    Input  packets:                   4
    Output packets:            15721796
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Ingress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    0                    0                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                   0                    0                    0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort            201985796            201985796                    0
```

```
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                  65                   65                    0

  Logical interface ae0.0 (Index 72) (SNMP ifIndex 505) (Generation 204)
    Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             7          0           440            0
        Output:      22768200     148466    1047338120     54635848
    Link:
      ge-2/1/6.0
        Input :             7          0           440            0
        Output:      22768200     148466    1047338120     54635848
    Marker Statistics:   Marker Rx     Resp Tx   Unknown Rx   Illegal Rx
      ge-2/1/6.0                0           0            0            0
    Protocol inet, MTU: 1500, Generation: 291, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.30.30.0/30, Local: 10.30.30.1, Broadcast: 10.30.30.3, Generation: 420
    Protocol inet6, MTU: 1500, Generation: 292, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: ::/26, Local: ::10.30.30.1
    Generation: 422
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::21f:12ff:fec2:37f0
    Protocol multiservice, MTU: Unlimited, Generation: 424
    Generation: 293, Route table: 0
      Policer: Input: __default_arp_policer__
```

show interfaces statistics (SONET/SDH)

```
user@host> show interfaces statistics detail so-3/0/0 | no-more
Physical interface: so-3/0/0, Enabled, Physical link is Up
  Interface index: 133, SNMP ifIndex: 538, Generation: 283
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC192, Loopback: None,
FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives
  Hold-times     : Up 0 ms, Down 0 ms
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
```

```
  Keepalive statistics:
    Input : 13 (last seen 00:00:04 ago)
    Output: 14 (last sent 00:00:02 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Opened, iso: Not-configured, mpls: Not-configured
  CHAP state: Closed
  PAP state: Closed
  CoS queues     : 8 supported, 8 maximum usable queues
  Last flapped   : 2009-11-09 02:52:34 PST (01:12:39 ago)
  Statistics last cleared: 2009-11-09 03:58:54 PST (00:06:19 ago)
  Traffic statistics:
   Input  bytes  :           2559160294             54761720 bps
   Output bytes  :                10640                   48 bps
   Input  packets:             55633975               148809 pps
   Output packets:                  216                    0 pps
   IPv6 transit statistics:
    Input  bytes  :           647922328
    Output bytes  :                   0
    Input  packets:            14085269
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Bucket drops: 0, Policed
discards: 0, L3 incompletes: 0,
    L2 channel errors: 0, L2 mismatch timeouts: 0, HS link CRC errors: 0, HS link FIFO
overflows: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Aged packets: 0, HS link FIFO underflows: 0,
MTU errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                  4                   4                   0
    1 expedited-fo                 0                   0                   0
    2 assured-forw                 0                   0                   0
    3 network-cont               213                 213                   0
  SONET alarms   : None
  SONET defects  : None

  Logical interface so-3/0/0.0 (Index 72) (SNMP ifIndex 578) (Generation 182)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
    Protocol inet, MTU: 4470, Generation: 244, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.30.30.0/30, Local: 10.30.30.2, Broadcast: 10.30.30.3, Generation: 322
    Protocol inet6, MTU: 4470, Generation: 245, Route table: 0
```

```
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: ::10.30.30.0/126, Local: ::10.30.30.2
    Generation: 324
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::2a0:a5ff:fe61:9264
    Generation: 326
```

show interfaces statistics (Aggregated SONET/SDH—Ingress)

```
user@host> show interfaces statistics detail as0 | no-more
Physical interface: as0, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 534, Generation: 282
  Link-level type: PPP, MTU: 4474, Speed: OC192, Minimum links needed: 1, Minimum bandwidth
needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Last flapped   : 2009-11-09 03:45:53 PST (00:09:38 ago)
  Statistics last cleared: 2009-11-09 03:48:17 PST (00:07:14 ago)
  Traffic statistics:
   Input  bytes  :           2969786332             54761688 bps
   Output bytes  :                11601                    0 bps
   Input  packets:             64560636               148808 pps
   Output packets:                  225                    0 pps
   IPv6 transit statistics:
    Input  bytes  :          2086013152
    Output bytes  :                   0
    Input  packets:            45348114
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort                    3                    3                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                 222                  222                    0
```

```
  Logical interface as0.0 (Index 71) (SNMP ifIndex 576) (Generation 179)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :      64560550     148808    2969785300     54761688
        Output:           139          0         10344            0
    Link:
      so-3/0/0.0
        Input :      64560550     148808    2969785300     54761688
        Output:           139          0         10344            0
    Protocol inet, MTU: 4470, Generation: 240, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.30.30.0/30, Local: 10.30.30.2, Broadcast: 10.30.30.3, Generation: 316
    Protocol inet6, MTU: 4470, Generation: 241, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: ::10.30.30.0/126, Local: ::10.30.30.2
    Generation: 318
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::2a0:a5ff:fe61:9264
    Generation: 320
```

show interfaces statistics (Aggregated SONET/SDH—Egress)

```
user@host> show interfaces statistics detail as0 | no-more
Physical interface: as0, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 565, Generation: 323
  Link-level type: PPP, MTU: 4474, Speed: OC192, Minimum links needed: 1, Minimum bandwidth
needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Last flapped   : 2009-11-09 03:43:37 PST (00:12:48 ago)
  Statistics last cleared: 2009-11-09 03:48:54 PST (00:07:31 ago)
  Traffic statistics:
   Input  bytes  :               11198                 392 bps
   Output bytes  :          3101452132            54783448 bps
   Input  packets:                 234                   0 pps
   Output packets:            67422937              148868 pps
   IPv6 transit statistics:
```

```
    Input  bytes  :                5780
    Output bytes  :          2171015678
    Input  packets:                  72
    Output packets:            47195993
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 best-effort             67422830             67422830                    0
    1 expedited-fo                   0                    0                    0
    2 assured-forw                   0                    0                    0
    3 network-cont                  90                   90                    0

  Logical interface as0.0 (Index 71) (SNMP ifIndex 548) (Generation 206)
    Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :           144          0         10118          392
        Output:      67422847     148868    3101450962     54783448
    Link:
      so-0/1/0.0
        Input :           144          0         10118          392
        Output:      67422847     148868    3101450962     54783448
    Protocol inet, MTU: 4470, Generation: 295, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.30.30.0/30, Local: 10.30.30.1, Broadcast: 10.30.30.3, Generation: 426
    Protocol inet6, MTU: 4470, Generation: 296, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: ::/26, Local: ::10.30.30.1
    Generation: 428
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::2a0:a5ff:fe63:1d0a
    Generation: 429
```

show interfaces statistics (MX Series Routers)

```
user@host> show interfaces xe-0/0/0 statistics
Physical interface: xe-0/0/0, Enabled, Physical link is Up
```

```
  Interface index: 145, SNMP ifIndex: 592
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error: None, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Pad to minimum frame size: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:00:5E:00:53:f0, Hardware address: 00:00:5E:00:53:f0
  Last flapped   : 2013-10-26 03:20:40 test (2w3d 03:29 ago)
  Statistics last cleared: Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0
  Active alarms  : LINK
  Active defects : LINK
  PCS statistics                      Seconds
    Bit errors                           109
    Errored blocks                       109
  Interface transmit statistics: Disabled
```

show interfaces statistics (MX Series Routers; With RPF Check Detail of Ethernet Interface)

```
user@host> show interfaces ge-0/0/0 extensive
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 149, SNMP ifIndex: 527, Generation: 152
  Link-level type: Ethernet, MTU: 1514, MRU: 1522, LAN-PHY mode, Speed: 1000mbps, BPDU Error:
None, Loop Detect PDU Error: None,
  Ethernet-Switching Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Pad to minimum frame size: Disabled
  Device flags   : Present Running
  Interface Specific flags: RPF_statistics
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Damping        : half-life: 0 sec, max-suppress: 0 sec, reuse: 0, suppress: 0, state:
unsuppressed
  Current address: 56:68:ad:c0:1f:72, Hardware address: 56:68:ad:c0:1f:72
```

```
  Last flapped   : 2021-02-12 17:21:37 IST (3d 15:55 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
IPv6 transit statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  Dropped traffic statistics due to STP State:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3 incompletes: 0, L2
channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0,
HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
 RPF counters statistics:
   Rpf IPv4 bytes   :                     0
   Rpf IPv4 packets :                   0
   Rpf IPv6 bytes   :                     0
   Rpf IPv6 packets :                   0
  Active alarms  : None
  Active defects : None
  PCS statistics                     Seconds
    Bit errors                            0
    Errored blocks                        0
```

show interfaces statistics (MX Series Routers: Dynamic Interfaces with RPF Check Detail)

```
user@host> show interfaces statistics pp0.3221225475 detail
  Logical interface pp0.3221225475(Index 536870921)(SNMP ifIndex 200000009) (Generation 6)
    Flags: Up Point-To-Point Encapsulation: PPPoE
    PPPoE:
```

```
      State: SessionUp, Session ID: 1,
      Session AC name: B, Remote MAC address:00:00:5E:00:53:01,
      Underlying interface: xe-1/0/0.3221225474 (Index 536870919)
      Ignore End-Of-List tag: Disable
    Bandwidth: 0
    Traffic statistics:
     Input  bytes  :                  34
     Output bytes  :                   0
     Input  packets:                   1
     Output packets:                   1
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                  34                   0 bps
     Output bytes  :                   0                   0 bps
     Input  packets:                   1                   0 pps
     Output packets:                   1                   0 pps
  Keepalive settings: Interval 30 seconds, Up-count 3, Down-count 3
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls: Not-configured
  CHAP state: Success
  PAP state: Closed
    Protocol inet, MTU: 1492
    Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0, NH drop cnt: 0
    Generation: 0, Route table: 0
      Flags: uRPF, Unnumbered
      RPF Failures: Packets: 0, Bytes: 0
      Donor interface: lo0.0 (Index 320)
      Input Filters: upstrm1-inet-pp0.3221225475-in
      Output Filters: dwnstrm1-inet-pp0.3221225475-out
      Addresses, Flags: Is-Primary
        Destination: Unspecified, Local: 10.255.96.19, Broadcast: Unspecified, Generation: 0
```

show interfaces statistics (PTX Series Packet Transport Routers)

```
user@host> show interfaces statistics em0
 Physical interface: em0, Enabled, Physical link is Up
      Interface index: 8, SNMP ifIndex: 0
```

```
      Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps
      Device flags   : Present Running
      Interface flags: SNMP-Traps
      Link type      : Full-Duplex
      Current address: 00:00:5E:00:53:1b, Hardware address: 00:00:5E:00:53:1b
      Last flapped   : Never
      Statistics last cleared: Never
    Input packets : 212620
    Output packets: 71
      Input errors: 0, Output errors: 0

      Logical interface em0.0 (Index 3) (SNMP ifIndex 0)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 212590
    Output packets: 71
    Protocol inet, MTU: 1500
      Flags: Is-Primary
      Addresses, Flags: Is-Default Is-Preferred Is-Primary
        Destination: 192.168.3/24, Local: 192.168.3.30,
        Broadcast: 192.168.3.255
```

show interfaces statistics (ACX Series routers)

```
user@host> show interfaces statistics ge-0/1/7
Physical interface: ge-0/1/7, Enabled, Physical link is Down
  Interface index: 151, SNMP ifIndex: 524
  Link-level type: Ethernet, Media type: Copper, MTU: 1514, Link-mode: Full-duplex, Speed:
1000mbps, BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled, Remote fault:
Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:00:5E:00:53:a3, Hardware address: 00:00:5E:00:53:a3
  Last flapped   : 2012-05-11 04:25:28 PDT (2d 20:23 ago)
  Statistics last cleared: 2012-05-13 23:07:23 PDT (01:41:25 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Input errors: 0, Output errors: 0
  Active alarms  : LINK
```

```
  Active defects : LINK
  Interface transmit statistics: Disabled
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 12.2 for ACX Series Routers.

`satellite-device` option introduced in Junos OS Release 14.2R3.

RELATED DOCUMENTATION

clear interfaces statistics

show passive-monitoring error

IN THIS SECTION

- Syntax | **423**
- Description | **423**
- Options | **423**
- Required Privilege Level | **423**
- Output Fields | **423**
- Sample Output | **425**
- Release Information | **426**

Syntax

```
show passive-monitoring error (* | all | mo-fpc/pic/port)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display passive monitoring error statistics.

Options

*** | all | mo-*fpc/pic/port*** Display error statistics for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name.

Required Privilege Level

view

Output Fields

Table 11 on page 423 lists the output fields for the `show passive-monitoring error` command. Output fields are listed in the approximate order in which they appear.

Table 11: show passive-monitoring error Output Fields

Field Name	Field Description
Passive monitoring interface	Name of the passive monitoring interface.

Table 11: show passive-monitoring error Output Fields *(Continued)*

Field Name	Field Description
Local interface index	Index counter of the local interface.
Interface state	State of the passive monitoring interface: • **Monitoring**—Specified interface is actively monitoring. • **Disabled**—Specified interface has been disabled from the CLI. • **Not monitoring**—The interface is operational, but not monitoring. This condition occurs when an interface first comes online, or when the interface is operational, but no logical unit has been configured under the physical interface. • **Unknown**—Unknown state. • **Error**—An error occurred during the process of determining the state of the interface.
Error information	
Packets dropped (no memory)	Number of packets dropped because of memory shortage.
Packets dropped (not IP)	Number of non-IP packets dropped.
Packets dropped (not IPv4)	Number of packets dropped because they failed the IPv4 version check.
Packets dropped (header too small)	Number of packets dropped because the packet length or IP header length was too small.
Memory allocation failures	Number of flow record memory allocation failures. A small number reflects failures to replenish the free list. A large number indicates the monitoring station is almost out of memory space.
Memory free failures	Number of flow record memory free failures.

Table 11: show passive-monitoring error Output Fields *(Continued)*

Field Name	Field Description
Memory free list failures	Number of flow records received from free list that failed. Memory is nearly exhausted or too many new flows greater than 128 KB are being created per second.
Memory warning	Whether the flows have exceeded 1 million packets per second (Mpps) on a Monitoring Services PIC or 2 Mpps on a Monitoring Services II PIC. The response can be **Yes** or **No**.
Memory overload	Whether the memory has been overloaded. The response can be **Yes** or **No.**
PPS overload	Whether the PIC is receiving more packets per second than the configured threshold. The response can be **Yes** or **No**.
BPS overload	Whether the PIC is receiving more bits per second than the configured threshold. The response can be **Yes** or **No**.

Sample Output

show passive-monitoring error all

```
user@host> show passive-monitoring error all
Passive monitoring interface: mo-4/0/0, Local interface index: 44
Interface state: Monitoring
  Error information
    Packets dropped (no memory): 0, Packets dropped (not IP): 0
    Packets dropped (not IPv4): 0, Packets dropped (header too small): 0
    Memory allocation failures: 0, Memory free failures: 0
    Memory free list failures: 0
    Memory warning: No, Memory overload: No, PPS overload: No, BPS overload: No

Passive monitoring interface: mo-4/1/0, Local interface index: 45
Interface state: Not monitoring
  Error information
    Packets dropped (no memory): 0, Packets dropped (not IP): 0
    Packets dropped (not IPv4): 0, Packets dropped (header too small): 0
```

```
    Memory allocation failures: 0, Memory free failures: 0
    Memory free list failures: 0
    Memory warning: No, Memory overload: No, PPS overload: No, BPS overload: No
```

Release Information

Command introduced before Junos OS Release 7.4.

show passive-monitoring flow

IN THIS SECTION

- Syntax | **426**
- Description | **426**
- Options | **427**
- Required Privilege Level | **427**
- Output Fields | **427**
- Sample Output | **429**
- Release Information | **429**

Syntax

```
show passive-monitoring flow (* | all | mo-fpc/pic/port)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display passive flow statistics.

Options

*** | all | mo-*fpc/pic/port*** Display passive flow statistics for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name.

Required Privilege Level

view

Output Fields

Table 12 on page 427 lists the output fields for the `show passive-monitoring flow` command. Output fields are listed in the approximate order in which they appear.

Table 12: show passive-monitoring flow Output Fields

Field Name	Field Description
Passive monitoring interface	Name of the passive monitoring interface.
Local interface index	Index counter of the local interface.
Interface state	State of the passive monitoring interface: • **Monitoring**—Specified interface is actively monitoring. • **Disabled**—Specified interface has been disabled from the CLI. • **Not monitoring**—The interface is operational, but not monitoring. This condition occurs when an interface first comes online, or when the interface is operational, but no logical unit has been configured under the physical interface. • **Unknown**—Unknown state. • **Error**—An error occurred during the process of determining the state of the interface.

Table 12: show passive-monitoring flow Output Fields *(Continued)*

Field Name	Field Description
Flow information	
Flow packets	Number of packets received by an operational PIC.
Flow bytes	Number of bytes received by an operational PIC.
Flow packets 10-second rate	Number of packets per second handled by the PIC and displayed as a 10-second average.
Flow bytes 10-second rate	Number of bytes per second handled by the PIC and displayed as a 10-second average.
Active flows	Number of currently active flows tracked by the PIC.
Total flows	Total number of flows received by an operational PIC.
Flows exported	Total number of flows exported by an operational PIC.
Flows packets exported	Total number of cflowd packets exported by an operational PIC.
Flows inactive timed out	Total number of flows that are exported because of inactivity.
Flows active timed out	Total number of long-lived flows that are exported because of an active timeout.

Sample Output

show passive-monitoring flow all

```
user@host> show passive-monitoring flow all
Passive monitoring interface: mo-4/0/0, Local interface index: 44
Interface state: Monitoring
  Flow information
    Flow packets: 6533434, Flow bytes: 653343400
    Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
    Active flows: 0, Total flows: 1599
    Flows exported: 1599, Flows packets exported: 55
    Flows inactive timed out: 1599, Flows active timed out: 0

Passive monitoring interface: mo-4/1/0, Local interface index: 45
Interface state: Monitoring
  Flow information
    Flow packets: 6537780, Flow bytes: 653778000
    Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
    Active flows: 0, Total flows: 1601
    Flows exported: 1601, Flows packets exported: 55
    Flows inactive timed out: 1601, Flows active timed out: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

show passive-monitoring memory

IN THIS SECTION

- Syntax | 430
- Description | 430

- Options | 430
- Required Privilege Level | 430
- Output Fields | 431
- Sample Output | 432
- Release Information | 432

Syntax

```
show passive-monitoring memory (* | all | mo-fpc/pic/port)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display passive monitoring memory and flow record statistics

Options

*** | all | mo-*fpc/pic/port*** Display memory and flow record statistics for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name.

Required Privilege Level

view

Output Fields

Table 13 on page 431 lists the output fields for the `show passive-monitoring memory` command. Output fields are listed in the approximate order in which they appear.

Table 13: show passive-monitoring memory Output Fields

Field Name	Field Description
Passive monitoring interface	Name of the passive monitoring interface.
Local interface index	Index counter of the local interface.
Memory utilization	
Allocation count	Number of flow records allocated.
Free count	Number of flow records freed.
Maximum allocated	Maximum number of flow records allocated since the monitoring station booted. This number represents the peak number of flow records allocated at a time.
Allocations per second	Flow records allocated per second during the last statistics interval on the PIC.
Frees per second	Flow records freed per second during the last statistics interval on the PIC.
Total memory used, Total memory free	Total memory currently used and total amount of memory currently free (in bytes).

Sample Output

show passive-monitoring memory all

```
user@host> show passive-monitoring memory all
Passive monitoring interface: mo-4/0/0, Local interface index: 44
  Memory utilization
    Allocation count: 1600, Free count: 1599, Maximum allocated: 1600
    Allocations per second: 3200, Frees per second: 1438
    Total memory used (in bytes): 103579176, Total memory free (in bytes):     163914184
```

Release Information

Command introduced before Junos OS Release 7.4.

show passive-monitoring status

IN THIS SECTION

- Syntax | 433
- Description | 433
- Options | 433
- Required Privilege Level | 433
- Output Fields | 433
- Sample Output | 434
- Release Information | 435

Syntax

```
show passive-monitoring status (*| all | mo-fpc/pic/port)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display passive monitoring status.

Options

*** | all | mo-*fpc/pic/port*** Display status for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name.

Required Privilege Level

view

Output Fields

Table 14 on page 433 lists the output fields for the `show passive-monitoring status` command. Output fields are listed in the approximate order in which they appear.

Table 14: show passive-monitoring status Output Fields

Output Field	Output Field Description
Passive monitoring interface	Name of the passive monitoring interface.
Local interface index	Index counter of the local interface.

Table 14: show passive-monitoring status Output Fields *(Continued)*

Output Field	Output Field Description
Interface state	Monitoring state of the passive monitoring interface. • **Monitoring**—PIC is actively monitoring. • **Disabled**—PIC has been disabled using the CLI. • **Not monitoring**—PIC is operational, but not monitoring. This condition can happen while the PIC is coming online, or when the PIC is operational but has no logical unit configured under the physical interface. • **Unknown**
Group index	Integer that represents the monitoring group of which the PIC is a member. **Group index** is a mapping from the group name to an index. It is not related to the number of monitoring groups.
Export interval	Configured export interval for cflowd records, in seconds.
Export format	Configured export format (only cflowd version 5 is supported).
Protocol	Protocol the PIC is configured to monitor (only IPv4 is supported).
Engine type	Configured engine type that is inserted in output cflowd packets.
Engine ID	Configured engine ID that is inserted in output cflowd packets.

Sample Output

show passive-monitoring status all

```
user@host> show passive-monitoring status all
Passive monitoring interface: mo-4/0/0, Local interface index: 44
Interface state: Monitoring
```

```
  Group index: 0
  Export interval: 15 secs, Export format: cflowd v5
  Protocol: IPv4, Engine type: 1, Engine ID: 1

Passive monitoring interface: mo-4/1/0, Local interface index: 45
Interface state: Disabled

Passive monitoring interface: mo-4/2/0, Local interface index: 46
Interface state: Not monitoring
```

Release Information

Command introduced before Junos OS Release 7.4.

show passive-monitoring usage

IN THIS SECTION

- Syntax | 435
- Description | 436
- Options | 436
- Required Privilege Level | 436
- Output Fields | 436
- Sample Output | 437
- Release Information | 437

Syntax

```
show passive-monitoring usage (* | all | mo-fpc/pic/port)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display passive monitoring usage statistics.

Options

*** | all | mo-*fpc/pic/port*** Display usage statistics for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name.

Required Privilege Level

view

Output Fields

Table 15 on page 436 lists the output fields for the `show passive-monitoring usage` command. Output fields are listed in the approximate order in which they appear.

Table 15: show passive-monitoring usage Output Fields

Output Field	Output Field Description
Passive monitoring interface	Name of the passive monitoring interface.
Local interface index	Index counter of the local interface.
CPU utilization	
Uptime	Time, in milliseconds, that the PIC has been operational.

Table 15: show passive-monitoring usage Output Fields *(Continued)*

Output Field	Output Field Description
Interrupt time	Total time that the PIC has spent processing packets since the last PIC reset.
Load (5 second)	CPU load on the PIC, averaged more than 5 seconds. The number is a percentage obtained by dividing the time spent on active tasks by the total elapsed time.
Load (1 minute)	CPU load on the PIC, averaged more than 1 minute. The number is a percentage obtained by dividing the time spent on active tasks by the total elapsed time.

Sample Output

show passive-monitoring usage all

```
user@host> show passive-monitoring usage
Passive monitoring interface: mo-4/0/0, Local interface index: 44
  CPU utilization
    Uptime: 653155 milliseconds, Interrupt time: 40213754 microseconds
    Load (5 second): 20%, Load (1 minute): 17%

Passive monitoring interface: mo-4/1/0, Local interface index: 45
  CPU utilization
    Uptime: 652292 milliseconds, Interrupt time: 40223178 microseconds
    Load (5 second): 22%, Load (1 minute): 15%

Passive monitoring interface: mo-4/2/0, Local interface index: 46
  CPU utilization
    Uptime: 649491 milliseconds, Interrupt time: 40173645 microseconds
    Load (5 second): 22%, Load (1 minute): 10098862%
```

Release Information

Command introduced before Junos OS Release 7.4.

show route forwarding-table

IN THIS SECTION

- Syntax | 438
- Syntax (MX Series Routers) | 439
- Syntax (TX Matrix and TX Matrix Plus Routers) | 439
- Description | 440
- Options | 440
- Required Privilege Level | 442
- Output Fields | 442
- Sample Output | 449
- Release Information | 452

Syntax

```
show route forwarding-table
<detail | extensive | summary>
<all>
<ccc interface-name>
<destination destination-prefix>
<family family | matching matching>
<interface-name interface-name>
<label name>
<matching matching>
<multicast>
<table (default | logical-system-name/routing-instance-name | routing-instance-name)>
<vlan (all | vlan-name)>
<vpn vpn>
```

Syntax (MX Series Routers)

```
show route forwarding-table
<detail | extensive | summary>
<all>
<bridge-domain (all | domain-name)>
<ccc interface-name>
<destination destination-prefix>
<family family | matching matching>
<interface-name interface-name>
<label name>
<learning-vlan-id learning-vlan-id>
<matching matching>
<multicast>
<table (default | logical-system-name/routing-instance-name | routing-instance-name)>
<vlan (all |  vlan-name)>
<vpn vpn>
```

Syntax (TX Matrix and TX Matrix Plus Routers)

```
show route forwarding-table
<detail | extensive | summary>
<all>
<ccc interface-name>
<destination destination-prefix>
<family family | matching matching>
<interface-name interface-name>
<matching matching>
<label name>
<lcc number>
<multicast>
<table routing-instance-name>
<vpn vpn>
```

Description

Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.

NOTE: The Routing Engine copies the forwarding table to the Packet Forwarding Engine, the part of the router that is responsible for forwarding packets. To display the entries in the Packet Forwarding Engine's forwarding table, use the `show pfe route` command.

Options

none	Display the routes in the forwarding tables. By default, the `show route forwarding-table` command does not display information about private, or internal, forwarding tables.
detail | extensive | summary	(Optional) Display the specified level of output.
all	(Optional) Display routing table entries for all forwarding tables, including private, or internal, tables.
bridge-domain (all | *bridge-domain-name*)	(MX Series routers only) (Optional) Display route entries for all bridge domains or the specified bridge domain.
ccc *interface-name*	(Optional) Display route entries for the specified circuit cross-connect interface.
destination *destination-prefix*	(Optional) Destination prefix.
family *family*	(Optional) Display routing table entries for the specified family: `bridge` (`ccc` | `destination` | `detail` | `extensive` | `interface-name` | `label` | `learning-vlan-id` | `matching` | `multicast` | `summary` | `table` | `vlan` | `vpn`), `ethernet-switching`, `evpn`, `fibre-channel`, `fmembers`, `inet`, `inet6`, `iso`, `mcsnoop-inet`, `mcsnoop-inet6`, `mpls`, `satellite-inet`, `satellite-inet6`, `satellite-vpls`, `tnp`, `unix`, `vpls`, or `vlan-classification`.
interface-name *interface-name*	(Optional) Display routing table entries for the specified interface.

label *name*

(Optional) Display route entries for the specified label.

lcc *number*

(TX Matrix and TX matrix Plus routers only) (Optional) On a routing matrix composed of a TX Matrix router and T640 routers, display information for the specified T640 router (or line-card chassis) connected to the TX Matrix router. On a routing matrix composed of the TX Matrix Plus router and T1600 or T4000 routers, display information for the specified router (line-card chassis) connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

learning-vlan-id *learning-vlan-id*

(MX Series routers only) (Optional) Display learned information for all VLANs or for the specified VLAN.

matching *matching*

(Optional) Display routing table entries matching the specified prefix or prefix length.

multicast

(Optional) Display routing table entries for multicast routes.

table

(Optional) Display route entries for all the routing tables in the main routing instance or for the specified routing instance. If your device supports logical systems, you can also display route entries for the specified logical system and routing instance. To view the routing instances on your device, use the `show route instance` command.

vlan (all | *vlan-name***)**

(Optional) Display information for all VLANs or for the specified VLAN.

vpn *vpn*

(Optional) Display routing table entries for a specified VPN.

Required Privilege Level

view

Output Fields

Table 16 on page 442 lists the output fields for the `show route forwarding-table` command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified, or when the `detail` keyword is used instead of the `extensive` keyword.

Table 16: show route forwarding-table Output Fields

Field Name	Field Description	Level of Output
Logical system	Name of the logical system. This field is displayed if you specify the `table` *`logical-system-name/routing-instance-name`* option on a device that is configured for and supports logical systems.	All levels
Routing table	Name of the routing table (for example, inet, inet6, mpls).	All levels

Table 16: show route forwarding-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Enabled protocols	The features and protocols that have been enabled for a given routing table. This field can contain the following values: • BUM hashing—BUM hashing is enabled. • MAC Stats—Mac Statistics is enabled. • Bridging—Routing instance is a normal layer 2 bridge. • No VLAN—No VLANs are associated with the bridge domain. • All VLANs—The `vlan-id all` statement has been enabled for this bridge domain. • Single VLAN—Single VLAN ID is associated with the bridge domain. • MAC action drop—New MACs will be dropped when the MAC address limit is reached. • Dual VLAN—Dual VLAN tags are associated with the bridge domain • No local switching—No local switching is enabled for this routing instance.. • Learning disabled—Layer 2 learning is disabled for this routing instance. • MAC limit reached—The maximum number of MAC addresses that was configured for this routing instance has been reached. • VPLS—The VPLS protocol is enabled. • No IRB l2-copy—The no-irb-layer-2-copy feature is enabled for this routing instance. • ACKed by all peers—All peers have acknowledged this routing instance. • BUM Pruning—BUM pruning is enabled on the VPLS instance. • Def BD VXLAN—VXLAN is enabled for the default bridge domain. • EVPN—EVPN protocol is enabled for this routing instance.	All levels

Table 16: show route forwarding-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
	• Def BD OVSDB—Open vSwitch Database (OVSDB) is enabled on the default bridge domain. • Def BD Ingress replication—VXLAN ingress node replication is enabled on the default bridge domain. • L2 backhaul—Layer 2 backhaul is enabled. • FRR optimize—Fast reroute optimization • MAC pinning—MAC pinning is enabled for this bridge domain. • MAC Aging Timer—The MAC table aging time is set per routing instance. • EVPN VXLAN—This routing instance supports EVPN with VXLAN encapsulation. • PBBN—This routing instance is configured as a provider backbone bridged network. • PBN—This routing instance is configured as a provider bridge network. • ETREE—The ETREE protocol is enabled on this EVPN routing instance. • ARP/NDP suppression—EVPN ARP NDP suppression is enabled in this routing instance. • Def BD EVPN VXLAN—EVPN VXLAN is enabled for the default bridge domain. • MPLS control word—Control word is enabled for this MPLS routing instance.	
Address family	Address family (for example, `IP`, `IPv6`, `ISO`, `MPLS`, and `VPLS`).	All levels
Destination	Destination of the route.	`detail extensive`

Table 16: show route forwarding-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Route Type (Type)	How the route was placed into the forwarding table. When the `detail` keyword is used, the route type might be abbreviated (as shown in parentheses): • `cloned` (`clon`)—(TCP or multicast only) Cloned route. • `destination` (`dest`)—Remote addresses directly reachable through an interface. • `destination down` (`iddn`)—Destination route for which the interface is unreachable. • `interface cloned` (`ifcl`)—Cloned route for which the interface is unreachable. • `route down` (`ifdn`)—Interface route for which the interface is unreachable. • `ignore` (`ignr`)—Ignore this route. • `interface` (`intf`)—Installed as a result of configuring an interface. • `permanent` (`perm`)—Routes installed by the kernel when the routing table is initialized. • `user`—Routes installed by the routing protocol process or as a result of the configuration.	All levels
Route Reference (RtRef)	Number of routes to reference.	`detail extensive`

Table 16: show route forwarding-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Flags	Route type flags: • `none`—No flags are enabled. • `accounting`—Route has accounting enabled. • `cached`—Cache route. • `incoming-iface` *`interface-number`*—Check against incoming interface. • `prefix load balance`—Load balancing is enabled for this prefix. • `rt nh decoupled`—Route has been decoupled from the next hop to the destination. • `sent to PFE`—Route has been sent to the Packet Forwarding Engine. • `static`—Static route.	`extensive`
Next hop	IP address of the next hop to the destination. **NOTE**: For static routes that use point-to-point (P2P) outgoing interfaces, the next-hop address is not displayed in the output.	`detail extensive`

Table 16: show route forwarding-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Next hop Type (Type)	Next-hop type. When the `detail` keyword is used, the next-hop type might be abbreviated (as indicated in parentheses): • `broadcast (bcst)`—Broadcast. • `deny`—Deny. • `discard (dscd)` —Discard. • `hold`—Next hop is waiting to be resolved into a unicast or multicast type. • `indexed (idxd)`—Indexed next hop. • `indirect (indr)`—Indirect next hop. • `local (locl)`—Local address on an interface. • `routed multicast (mcrt)`—Regular multicast next hop. • `multicast (mcst)`—Wire multicast next hop (limited to the LAN). • `multicast discard (mdsc)`—Multicast discard. • `multicast group (mgrp)`—Multicast group member. • `receive (recv)`—Receive. • `reject (rjct)`—Discard. An ICMP unreachable message was sent. • `resolve (rslv)`—Resolving the next hop. • `unicast (ucst)`—Unicast. • `unilist (ulst)`—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list. • `VxLAN Local`—EVPN Type 5 route in kernel.	`detail extensive`
Index	Software index of the next hop that is used to route the traffic for a given prefix.	`detail extensive` none

Table 16: show route forwarding-table Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Route interface-index	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	`extensive`
Reference (NhRef)	Number of routes that refer to this next hop.	`detail extensive` none
Next-hop interface (Netif)	Interface used to reach the next hop.	`detail extensive` none
Weight	Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible (see the `Balance` field description).	`extensive`
Balance	Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a router is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.	`extensive`
RPF interface	List of interfaces from which the prefix can be accepted. Reverse path forwarding (RPF) information is displayed only when `rpf-check` is configured on the interface.	`extensive`
sftw	On QFX5130 Series devices that run Junos OS Evolved, you see the next-hop type as `sftw` in the output of the `show route forwarding-table` command. A software next-hop can contain a link to another single next-hop, and hence it often works as an indirect next-hop. The usage of `sftw` helps to reduce churn in Junos OS Evolved objects by providing a common object where topology updates can be reflected. Further MPLS states can also be attached to these software next-hops. When a route is created, a `sftw` next-hop seems to be created whether it is the management port or the loopback interface.	`detail extensive` none

Sample Output

show route forwarding-table

```
user@host> show route forwarding-table
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    46     4
0.0.0.0/32         perm     0                    dscd    44     1
172.16.1.0/24         ifdn     0                    rslv   608     1 ge-2/0/1.0
172.16.1.0/32         iddn     0 172.16.1.0            recv   606     1 ge-2/0/1.0
172.16.1.1/32         user     0                    rjct    46     4
172.16.1.1/32         intf     0 172.16.1.1            locl   607     2
172.16.1.1/32         iddn     0 172.16.1.1            locl   607     2
172.16.1.255/32       iddn     0 ff:ff:ff:ff:ff:ff  bcst   605     1 ge-2/0/1.0
10.0.0.0/24        intf     0                    rslv   616     1 ge-2/0/0.0
10.0.0.0/32        dest     0 10.0.0.0           recv   614     1 ge-2/0/0.0
10.0.0.1/32        intf     0 10.0.0.1           locl   615     2
10.0.0.1/32        dest     0 10.0.0.1           locl   615     2
10.0.0.255/32      dest     0 10.0.0.255         bcst   613     1 ge-2/0/0.0
10.1.1.0/24        ifdn     0                    rslv   612     1 ge-2/0/1.0
10.1.1.0/32        iddn     0 10.1.1.0           recv   610     1 ge-2/0/1.0
10.1.1.1/32        user     0                    rjct    46     4
10.1.1.1/32        intf     0 10.1.1.1           locl   611     2
10.1.1.1/32        iddn     0 10.1.1.1           locl   611     2
10.1.1.255/32      iddn     0 ff:ff:ff:ff:ff:ff  bcst   609     1 ge-2/0/1.0
10.206.0.0/16      user     0 10.209.63.254      ucst   419    20 fxp0.0
10.209.0.0/16      user     1 0:12:1e:ca:98:0    ucst   419    20 fxp0.0
10.209.0.0/18      intf     0                    rslv   418     1 fxp0.0
10.209.0.0/32      dest     0 10.209.0.0         recv   416     1 fxp0.0
10.209.2.131/32    intf     0 10.209.2.131       locl   417     2
10.209.2.131/32    dest     0 10.209.2.131       locl   417     2
10.209.17.55/32    dest     0 0:30:48:5b:78:d2   ucst   435     1 fxp0.0
10.209.63.42/32    dest     0 0:23:7d:58:92:ca   ucst   434     1 fxp0.0
10.209.63.254/32   dest     0 0:12:1e:ca:98:0    ucst   419    20 fxp0.0
10.209.63.255/32   dest     0 10.209.63.255      bcst   415     1 fxp0.0
10.227.0.0/16      user     0 10.209.63.254      ucst   419    20 fxp0.0

...
```

```
Routing table: iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                   rjct    27     1
47.0005.80ff.f800.0000.0108.0003.0102.5524.5220.00
intf     0                   locl    28     1

Routing table: inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                   rjct     6     1
ff00::/8           perm     0                   mdsc     4     1
ff02::1/128        perm     0 ff02::1           mcst     3     1

Routing table: ccc
MPLS:
Interface.Label    Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct 16     1
100004(top)fe-0/0/1.0
```

show route forwarding-table detail

```
user@host> show route forwarding-table detail
Routing table: inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            user     2 0:90:69:8e:b1:1b   ucst   132     4 fxp0.0
default            perm     0                    rjct    14     1
10.1.1.0/24        intf     0 ff.3.0.21          ucst   322     1 so-5/3/0.0
10.1.1.0/32        dest     0 10.1.1.0           recv   324     1 so-5/3/0.0
10.1.1.1/32        intf     0 10.1.1.1           locl   321     1
10.1.1.255/32      dest     0 10.1.1.255         bcst   323     1 so-5/3/0.0
10.21.21.0/24      intf     0 ff.3.0.21          ucst   326     1 so-5/3/0.0
10.21.21.0/32      dest     0 10.21.21.0         recv   328     1 so-5/3/0.0
10.21.21.1/32      intf     0 10.21.21.1         locl   325     1
10.21.21.255/32    dest     0 10.21.21.255       bcst   327     1 so-5/3/0.0
127.0.0.1/32       intf     0 127.0.0.1          locl   320     1
172.17.28.19/32    clon     1 192.168.4.254      ucst   132     4 fxp0.0
172.17.28.44/32    clon     1 192.168.4.254      ucst   132     4 fxp0.0

...
```

```
Routing table: private1__.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    46     1
10.0.0.0/8         intf     0                    rslv   136     1 fxp1.0
10.0.0.0/32        dest     0 10.0.0.0           recv   134     1 fxp1.0
10.0.0.4/32        intf     0 10.0.0.4           locl   135     2
10.0.0.4/32        dest     0 10.0.0.4           locl   135     2

...

Routing table: iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    38     1

Routing table: inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    22     1
ff00::/8           perm     0                    mdsc    21     1
ff02::1/128        perm     0 ff02::1            mcst    17     1

...

Routing table: mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct 28     1
```

show route forwarding-table destination extensive (EVPN Type 5 route with Type 2 and Type 5 route coexistence)

```
user@device> show route forwarding-table destination 10.1.1.20 table vrf1 extensive
Routing table: vrf1.inet [Index 9]
Internet:
Destination:  10.1.1.20/32
  Route type: user
  Route reference: 0                   Route interface-index: 0
  Multicast RPF nh index: 0
  P2mpidx: 0
```

```
  Flags: sent to PFE, VxLAN Local
  Nexthop:
  Next-hop type: composite             Index: 2694     Reference: 7
  Next-hop type: indirect              Index: 524326   Reference: 2
  Next-hop type: unilist               Index: 524288   Reference: 5
  Nexthop: 10.1.1.1
  Next-hop type: unicast               Index: 1724     Reference: 15
  Next-hop interface: xe-0/0/1.0    Weight: 0x0
  Nexthop: 10.1.1.4  Next-hop type: unicast               Index: 1725     Reference: 15
  Next-hop interface: xe-0/0/4.0    Weight: 0x0
```

show route forwarding-table extensive (RPF)

The next example is based on the following configuration, which enables an RPF check on all routes that are learned from this interface, including the interface route:

```
so-1/1/0 {
    unit 0 {
        family inet {
            rpf-check;
            address 192.0.2.2/30;
        }
    }
}
```

show route forwarding-table (QFX5130 Series Swiches)

```
user@QFX5130-32CD-r002> show route forwarding-table | grep sftw
default            user     0                    sftw     7021     1 re0:mgmt-0.0
200.0.0.2/32       user     0                    sftw     7022     1 et-0/0/32.0

user@jtac-QFX5130-32CD-r003> show route forwarding-table | grep sftw
default            user     0                    sftw     5021     1 re0:mgmt-0.0
200.0.0.1/32       user     0                    sftw     5022     1 et-0/0/32.0
```

Release Information

Command introduced before Junos OS Release 7.4.

Option `bridge-domain` introduced in Junos OS Release 7.5

Option `learning-vlan-id` introduced in Junos OS Release 8.4

Options `all` and `vlan` introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

show route instance

show services accounting aggregation

IN THIS SECTION

- Syntax | 453
- Description | 454
- Options | 454
- Additional Information | 455
- Required Privilege Level | 455
- Output Fields | 455
- Sample Output | 457
- Release Information | 459

Syntax

```
show services accounting aggregation aggregation-type <aggregation-value>
<detail | extensive | terse>
<limit limit-value>
 <  name service-name>
<order (bytes | packets)>
```

Description

Display information about the aggregated active flows being processed by the accounting service.

Options

aggregation-type <aggregation-value>

Display information for the specified aggregation type and optional value:

- **as** *<source-as-value* **|** *destination-as-value* **|** *input-snmp-interface-index-value* **|** *output-snmp-interface-index-value>*—Aggregate by autonomous system (AS).
- **destination-prefix <***destination-prefix-value* **|** *destination-as-value* **|** *output-snmp-interface-index-value>*—Aggregate by destination prefix.
- **protocol-port <***protocol-value* **|** *source-port-value* **|** *destination-port-value>*—Aggregate by protocol and port.
- **source-destination-prefix <***source-prefix-value* **|** *destination-prefix-value* **|** *destination-as-value* **|** *source-as-value* **|** *input-snmp-interface-index-value* **|** *output-snmp-interface-index-value>*—Aggregate by source and destination prefix.
- **source-prefix <***source-prefix-value* **|** *source-as-value* **|** *input-snmp-interface-index-value>*—Aggregate by source prefix.

detail | extensive | terse

(Optional) Display the specified level of output.

limit *limit-value*

(Optional) Limit the display output to the specified number of flows. The default is no limit.

name *service-name*

(Optional) Display information about the aggregated flows for a specified service name.

order (bytes | packets)

(Optional) Display the flow with the ordering of the highest number, either by byte count or by packet count.

Additional Information

For information about aggregation configuration options, see the Junos OS Services Interfaces Library for Routing Devices.

Required Privilege Level

view

Output Fields

Table 17 on page 455 lists the output fields for the `show services accounting aggregation` command. Output fields are listed in the approximate order in which they appear.

Table 17: show services accounting aggregation Output Fields

Field Name	Field Description
Service Accounting interface	Name of the service accounting interface.
Local interface index	Index corresponding to the service accounting interface.
Service name	Name of a service that was configured at the [`edit forwarding-options accounting`] hierarchy level. The default display, **(default sampling)**, indicates the service was configured at the [**edit forwarding-options sampling-level**] hierarchy level.
Protocol	Protocol identifier and number.
Source Port	Source port identifier and number.
Destination Port	Destination port identifier and number.
Source-AS	Source autonomous system (AS) number.

Table 17: show services accounting aggregation Output Fields *(Continued)*

Field Name	Field Description
Destination-AS	Destination AS number.
Source Prefix	Source prefix.
Destination Prefix	Destination prefix.
Source address	Source address.
Source prefix length	Source prefix length.
Destination address	Destination address.
Destination prefix length	Destination prefix length.
Input SNMP interface index	SNMP index of the interface the packet came in on.
Output SNMP interface index	SNMP index of the interface the packet went out on.
Start time	Actual time when the packet in this aggregation was first seen.
End time	Actual time when the packet in this aggregation was last seen.
Flow count	Number of flows in the aggregation.
Packet count	Number of packets in the aggregation.
Byte count	Number of bytes in the aggregation.

Sample Output

show services accounting aggregation protocol-port detail

```
user@host> show service accounting aggregation protocol-port detail
Service Accounting interface: mo-2/0/0, Local interface index: 468
Service name: (default sampling)
  Protocol: 6, Source port: 20, Destination port: 20
  Start time: 442349, End time: 6425714
  Flow count: 194, Packet count: 4294964388, Byte count: 4294781184

  Protocol: 0, Source port: 0, Destination port: 0
  Start time: 442349, End time: 6425749
  Flow count: 204, Packet count: 4294964324, Byte count: 4294777088

  Protocol: 17, Source port: 123, Destination port: 123
  Start time: 442364, End time: 6425784
  Flow count: 186, Packet count: 4294964152, Byte count: 4294766080
```

show services accounting aggregation source-destination-prefix

```
user@host> show service accounting aggregation source-destination-prefix
Service Accounting interface: rsp0, Local interface index: 171
Service name: (default sampling)
Interface state: Accounting
Source          Destination       Input        Output      Flow   Packet  Byte
prefix          prefix            interface    interface   count  count   count
192.0.2.0/20    198.51.100.0/24   ge-5/0/1.0   ge-5/0/0.0  256    491761  31472704
192.0.2.0/20    203.0.113.36/32   ge-5/0/1.0   ge-5/0/0.0  1      1926    123264
192.0.2.0/20    203.0.113.59/32   ge-5/0/1.0   ge-5/0/0.0  1      1926    123264
192.0.2.0/20    192.168.0.63/32   ge-5/0/1.0   ge-5/0/0.0  1      1925    123200
192.0.2.0/20    192.168.0.32/32   ge-5/0/1.0   ge-5/0/0.0  1      1925
```

show services accounting aggregation source-destination- prefix order packet detail

```
user@host> show service accounting aggregation source-destination-prefix order packet detail
name t2 input-snmp-interface-index 538
Service Accounting interface: mo-2/0/0, Local interface index: 468
```

```
Service name: t2
Source         Destination  Input SNMP  Output  SNMP   Flow   Packet Byte
Prefix            Prefix       Index      Index Count  Count    Count
10.1.1.2/20   192.168.167.1/0  538         432    1     60    46483
10.1.1.2/20   192.168.168.1/0  538         432    1     60     5191
10.1.1.2/20   192.168.154.1/0  538         432    2     60    45504
10.1.1.2/20   192.168.76.1/0   538         432    1     60    42177
10.1.1.2/20   192.168.149.1/0  538         432    1     60    49184
10.1.1.2/20   192.168.113.1/0  538         432    2     60    48757
```

show services accounting aggregation source-destination- prefix extensive limit

```
user@host> show service accounting aggregation source-destination-prefix name t2 extensive limit
3
Service Accounting interface: mo-2/0/0, Local interface index: 542
Service name: t2

  Source address: 10.1.1.2, Source prefix length: 20
  Destination address: 192.168.200.176.1, Destination prefix length: 0
  Input SNMP interface index: 24, Output SNMP interface index: 26
  Source-AS: 69, Destination-AS: 69
  Start time: Fri Feb 21 14:16:57 2003, End time: Fri Feb 21 14:22:50 2003
  Flow count: 0, Packet count: 6, Byte count: 5340

  Source address: 10.1.1.2, Source prefix length: 20
  Destination address: 192.168.160.1, Destination prefix length: 0
  Input SNMP interface index: 24, Output SNMP interface index: 26
  Source-AS: 69, Destination-AS: 69
  Start time: Fri Feb 21 14:16:57 2003, End time: Fri Feb 21 14:22:50 2003
  Flow count: 0, Packet count: 6, Byte count: 5490

  Source address: 10.1.1.2, Source prefix length: 20
  Destination address: 192.168.160.1, Destination prefix length: 0
  Input SNMP interface index: 24, Output SNMP interface index: 26
  Source-AS: 69, Destination-AS: 69
  Start time: Fri Feb 21 14:16:57 2003, End time: Fri Feb 21 14:22:50 2003
  Flow count: 0, Packet count: 6, Byte count: 4079
```

show services accounting aggregation source-destination-prefix name terse

```
user@host> show service accounting aggregation source-destination-prefix name T3 terse
Service Accounting interface: rsp0, Local interface index: 171
Service name: T3
Interface state: Accounting
Source        Destination       Input        Output       Flow   Packet  Byte
prefix        prefix            interface    interface    count  count   count
10.1.0.0/20   192.168.3.0/24    ge-5/0/1.0   ge-5/0/0.0   256    639822  40948608
10.1.0.0/20   192.168.2.67/32   ge-5/0/1.0   ge-5/0/0.0   1      2485    159040
10.1.0.0/20   192.168.2.92/32   ge-5/0/1.0   ge-5/0/0.0   1      2485
```

Release Information

Command introduced before Junos OS Release 7.4.

show services accounting aggregation template

IN THIS SECTION

- Syntax | **460**
- Description | **460**
- Options | **460**
- Required Privilege Level | **460**
- Output Fields | **460**
- Sample Output | **461**
- Release Information | **461**

Syntax

```
show services accounting aggregation template
<template-name template-name>
```

Description

Display information for flow aggregation version 9 templates.

Options

none	Display information for all flow aggregation version version 9 templates.
template-name *template-name*	(Optional) Display information for the specified template only.

Required Privilege Level

view

Output Fields

Table 18 on page 460 lists the output fields for the `show services accounting aggregation template` command. Output fields are listed in the approximate order in which they appear.

Table 18: show services accounting aggregation template Output Fields

Field Name	Field Description
MPLS Label 1	Position of first MPLS label.

Table 18: show services accounting aggregation template Output Fields *(Continued)*

Field Name	Field Description
MPLS Label 2	Position of second MPLS label.
MPLS Label 3	Position of third MPLS label.
MPLS Top Level Address	Outer top label FEC IP address.
Packet Count	Number of packets sent.

Sample Output

show services accounting aggregation template template-name

```
user@host> show services accounting aggregation template template-name mpls
MPLS label 1: 299808, MPLS label 2: 0, MPLS label 3: 0
Source address: 192.0.2.2, Destination address: 10.255.15.22, Top Label Address: 198.51.100.10
Source port: 0, Destination port: 0
Protocol: 61, TOS: 0, TCP flags: 0
Source mask: 24, Destination mask: 32
Input SNMP interface index: 503, Output SNMP interface index: 505
Start time: 40780, End time: 157330
Packet count: 3949198, Byte count: 181663062
```

Release Information

Command introduced in Junos OS Release 8.3.

show services accounting errors

IN THIS SECTION

- Syntax | 462
- Description | 462
- Options | 462
- Required Privilege Level | 463
- Output Fields | 463
- Sample Output | 465
- Sample Output | 465
- Release Information | 468

Syntax

```
show services accounting  errors
<inline-jflow | name (* | all | service-name)>
```

Description

Display active flow error statistics.

Options

none	Display error statistics for all services accounting instances.
inline-jflow fpc-slot *slot-number*	(Optional) Display error statistics for inline jflow.

name (* | all | *service-name*) (Optional) Display active flow error statistics. Use a wildcard character, specify all services, or provide a specific service name.

Required Privilege Level

view

Output Fields

Table 19 on page 463 lists the output fields for the `show services accounting errors` command. Output fields are listed in the approximate order in which they appear.

Table 19: show services accounting errors Output Fields

Field	Field Description
Service Accounting interface	Name of the service accounting interface.
Local interface index	Index counter of the local interface.
FPC slot	Slot number of the FPC for which the flow information is displayed. (Available only when the `inline-jflow fpc-slot` *`slot-number`* option is used.)
Service name	Name of a service that was configured at the [**edit forwarding-options accounting**] hierarchy level. The default display, **(default sampling)**, indicates the service was configured at the [**edit forwarding-options sampling-level**] hierarchy level.
Error Information	
Packets dropped (no memory)	Number of packets dropped because of memory shortage.

Table 19: show services accounting errors Output Fields *(Continued)*

Field	Field Description
Packets dropped (not IP)	Number of non-IP packets dropped.
Packets dropped (not IPv4)	Number of packets dropped because they failed the IPv4 version check.
Packets dropped (header too small)	Number of packets dropped because the packet length or IP header length was too small.
Memory allocation failures	Number of flow record memory allocation failures. A small number reflects failures to replenish the free list. A large number indicates the monitoring station is almost out of memory space.
Memory free failures	Number of flow record memory free failures.
Memory free list failures	Number of flow records received from the free list that failed. Memory is nearly exhausted, or too many new flows greater than 128 KB are being created per second.
Memory overload	Whether the memory has been overloaded. The response can be **Yes** or **No**.
PPS overload	Whether the PIC is receiving more packets per second than the configured threshold. The response can be **Yes** or **No**.
BPS overload	Whether the PIC is receiving more bits per second than the configured threshold. The response can be **Yes** or **No**.
Flow Creation Failures	Number of times flow creation failed.
Route Record Lookup Failures	Number of times the route record lookup failed.
AS Lookup Failures	Number of times autonomous system lookup failed.

Table 19: show services accounting errors Output Fields *(Continued)*

Field	Field Description
Export Packet Failures	Number of times packet export failed.

Sample Output

show services accounting errors (Monitoring PIC interface)

```
user@host> show services accounting errors
Service Accounting interface: mo-1/1/0, Local interface index: 15
Service name: (default sampling)
  Error information
    Packets dropped (no memory): 0, Packets dropped (not IP): 0
    Packets dropped (not IPv4): 0, Packets dropped (header too small): 0
    Memory allocation failures: 0, Memory free failures: 0
    Memory free list failures: 0
    Memory overload: No, PPS overload: No, BPS overload: No
```

Sample Output

show services accounting errors (Service PIC interface)

```
user@host> show services accounting errors
Service Accounting interface: sp-0/1/0
Service name: (default sampling)
  Error information
    Service sets dropped: 0, Active timeout failures: 0
    Export packet failures: 0, Flow creation failures: 0
    Memory overload: No
```

```
Service Accounting interface: sp-1/0/0
Service name: (default sampling)
  Error information
    Service sets dropped: 0, Active timeout failures: 0
    Export packet failures: 0, Flow creation failures: 0
    Memory overload: No
```

show services accounting errors inline-jflow fpc-slot (When Only IPv6 Is Configured)

```
user@host> show services accounting errors inline-jflow fpc-slot 5
  Error information
    FPC Slot: 5
    Flow Creation Failures: 0
    Route Record Lookup Failures: 0, AS Lookup Failures: 0
    Export Packet Failures: 0
    Memory Overload: No, Memory Alloc Fail Count: 0
```

show services accounting errors inline-jflow fpc-slot (When IPv4, IPv6, VPLS, and Bridge Are Configured)

```
user@host> show services accounting errors inline-jflow fpc-slot 5
  Error information
    FPC Slot: 5
    Flow Creation Failures: 0
    Route Record Lookup Failures: 0, AS Lookup Failures: 0
    Export Packet Failures: 0
    Memory Overload: No, Memory Alloc Fail Count: 0

    IPv4:
    IPv4 Flow Creation Failures: 0
    IPv4 Route Record Lookup Failures: 0, IPv4 AS Lookup Failures: 0
    IPv4 Export Packet Failures: 0

    IPv6:
    IPv6 Flow Creation Failures: 0
    IPv6 Route Record Lookup Failures: 0, IPv6 AS Lookup Failures: 0
    IPv6 Export Packet Failures: 0
```

```
    VPLS:
    VPLS Flow Creation Failures: 0
    VPLS Export Packet Failures: 0

    BRIDGE:
    BRIDGE Flow Creation Failures: 0
    BRIDGE Route Record Lookup Failures: 0, BRIDGE AS Lookup Failures: 0
    BRIDGE Export Packet Failures: 0
```

show services accounting errors inline-jflow (MX80 Router When Both IPv4 and IPv6 Are Configured)

```
user@host> show services accounting errors inline-jflow
  Error information
    TFEB Slot: 0
    Flow Creation Failures: 0
    Route Record Lookup Failures: 0, AS Lookup Failures: 0
    Export Packet Failures: 0
    Memory Overload: No

         IPv4:
    IPv4 Flow Creation Failures: 0
    IPv4 Route Record Lookup Failures: 0, IPv4 AS Lookup Failures: 0
    IPv4 Export Packet Failures: 0

    IPv6:
    IPv6 Flow Creation Failures: 0
    IPv6 Route Record Lookup Failures: 0, IPv6 AS Lookup Failures: 0
    IPv6 Export Packet Failures: 0
```

show services accounting errors inline-jflow fpc-slot(PTX1000 Router When Both IPv4 and IPv6 Are Configured)

```
user@host> show services accounting errors inline-jflow fpc-slot 0
Error information
FPC Slot: 0
Flow Creation Failures: 0
Route Record Lookup Failures: 0, AS Lookup Failures: 0
```

```
Export Packet Failures: 0
Memory Overload: No, Memory Alloc Fail Count: 0

IPv4:
IPv4 Flow Creation Failures: 0
IPv4 Route Record Lookup Failures: 0, IPv4 AS Lookup Failures: 0
IPv4 Export Packet Failures: 0

IPv6:
IPv6 Flow Creation Failures: 0
IPv6 Route Record Lookup Failures: 0, IPv6 AS Lookup Failures: 0
IPv6 Export Packet Failures: 0
```

show services accounting errors inline-jflow (SRX Series Devices When Both IPv4 and IPv6 Are Configured)

```
user@host> show services accounting errors inline-jflow
Error information
    FPC Slot: 0
    Flow Creation Failures: 0
    Route Record Lookup Failures: 0, AS Lookup Failures: 0
    Export Packet Failures: 0
    Memory Overload: No, Memory Alloc Fail Count: 0

    IPv4:
    IPv4 Flow Creation Failures: 0
    IPv4 Route Record Lookup Failures: 0, IPv4 AS Lookup Failures: 0
    IPv4 Export Packet Failures: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show services accounting flow

show services accounting flow

IN THIS SECTION

- Syntax | 469
- Description | 469
- Options | 469
- Required Privilege Level | 470
- Output Fields | 470
- Sample Output | 472
- Release Information | 478

Syntax

```
show services accounting flow
<inline-jflow fpc-slot slot-number | logical-system (all | logical-system) | name (* | all |
service-name)>
```

Description

Display active flow statistics.

Options

none	Display active flow statistics for all service instances.
logical-system (all | *logical-system*)	(Optional) Display active flow statistics for the specified logical system or all logical systems on the device.

inline-jflow (fpc-slot *slot-number*) (Optional) Display inline flow statistics for the specified FPC.

For PTX Series, starting in Junos OS Evolved Release 21.2R1 and Junos OS Release 21.3R1, every sampled packet is considered to be a flow. When the sampled packet is received, the flow is created and immediately timed out as inactive, and the software exports a record to the collector. Therefore, when you specify this option, the command now displays 0 for all of the various `Active Flows` and `Timed Out` fields. The values of the various `Total Flows` fields are now equal to their respective `Flow Packets` field values. The values of the various `Flows Inactive Timed Out` fields are now equal to their respective `Flow Packets` field values.

name (* | all | *service-name*) (Optional) Display services accounting active flow statistics. Use a wildcard character, specify all services, or provide a specific service name.

Required Privilege Level

view

Output Fields

Table 20 on page 470 lists the output fields for the `show services accounting flow` command. Output fields are listed in the approximate order in which they appear.

Table 20: show services accounting flow Output Fields

Output Field	Output Field Description
Service Accounting interface	Name of the service accounting interface.
Local interface index	Index counter of the local interface.
Service name	Name of a service that was configured at the [`edit forwarding-options accounting`] hierarchy level. The default display, **(default sampling)**, indicates the service was configured at the [**edit forwarding-options sampling-level**] hierarchy level.

Table 20: show services accounting flow Output Fields *(Continued)*

Output Field	Output Field Description
Flow Information	
FPC Slot	Slot number of the FPC for which the flow information is displayed. (Available only when the `inline-jflow fpc-slot` *slot-number* option is used.)
Flow packets	Number of packets received by an operational PIC.
Flow bytes	Number of bytes received by an operational PIC.
Flow packets 10-second rate	Number of packets per second handled by the PIC and displayed as a 10-second average.
Flow bytes 10-second rate	Number of bytes per second handled by the PIC and displayed as a 10-second average.
Active flows	Number of currently active flows tracked by the PIC.
Total flows	Total number of flows received by an operational PIC.
Flows exported	Total number of flows exported by an operational PIC.
Flows packets exported	Total number of cflowd packets exported by an operational PIC.
Flows inactive timed out	Total number of flows that are exported because of inactivity.
Flows active timed out	Total number of long-lived flows that are exported because of an active timeout.

Sample Output

show services accounting flow (Flow Aggregation v5/v8 Configuration)

```
user@host> show services accounting flow
Service Accounting interface: rsp0, Local interface index: 171
Service name: (default sampling)
Interface state: Accounting
  Flow information
    Flow packets: 87168293, Flow bytes: 5578770752
    Flow packets 10-second rate: 45762, Flow bytes 10-second rate: 2928962
    Active flows: 1000, Total flows: 2000
    Flows exported: 19960, Flows packets exported: 582
    Flows inactive timed out: 1000, Flows active timed out: 29000
```

show services accounting flow (Flow Aggregation v9 Configuration)

```
user@host> show services accounting flow
  Flow information
    Service Accounting interface: sp-7/1/0, Local interface index: 149
    Flow packets: 0, Flow bytes: 0
    Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
    Active flows: 0, Total flows: 0
    Flows exported: 0, Flows packets exported: 1
    Flows inactive timed out: 0, Flows active timed out: 0
```

show services accounting flow name

```
user@host> show services accounting flow name count2
Service Accounting interface: mo-1/1/0, Local interface index: 15
Service name: count2
  Flow information
    Flow packets: 0, Flow bytes: 0
    Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
    Active flows: 0, Total flows: 0
    Flows exported: 0, Flows packets exported: 0
    Flows inactive timed out: 0, Flows active timed out: 0
```

show services accounting flow name all

```
user@host> show services accounting flow name all
Service Accounting interface: rsp0, Local interface index: 171
Service name: T2
Interface state: Accounting
  Flow information
    Flow packets: 37609891, Flow bytes: 2407033024
    Flow packets 10-second rate: 45762, Flow bytes 10-second rate: 2928953
    Active flows: 1000, Total flows: 1000
    Flows exported: 6705, Flows packets exported: 198
    Flows inactive timed out: 0, Flows active timed out: 13000

Service Accounting interface: rsp0, Local interface index: 171
Service name: T3
Interface state: Accounting
  Flow information
    Flow packets: 37750807, Flow bytes: 2416051712
    Flow packets 10-second rate: 45762, Flow bytes 10-second rate: 2928940
    Active flows: 1000, Total flows: 1000
    Flows exported: 13437, Flows packets exported: 378
    Flows inactive timed out: 0, Flows active timed out: 13000

Service Accounting interface: rsp0, Local interface index: 171
Service name: T4
 Interface state: Accounting
  Flow information
    Flow packets: 0, Flow bytes: 0
    Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
    Active flows: 0, Total flows: 0
    Flows exported: 0, Flows packets exported: 0
    Flows inactive timed out: 0, Flows active timed out: 0

Service Accounting interface: rsp0, Local interface index: 171
Service name: count1
Interface state: Accounting
  Flow information
    Flow packets: 0, Flow bytes: 0
    Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
    Active flows: 0, Total flows: 0
```

```
    Flows exported: 0, Flows packets exported: 0
    Flows inactive timed out: 0, Flows active timed out: 0
```

show services accounting flow (Multiple Sampling Instances)

```
user@host> show services accounting flow
  Flow information
    Service Accounting interface: sp-2/0/0, Local interface index: 215
    Flow packets: 9867, Flow bytes: 631488
    Flow packets 10-second rate: 0, Flow bytes 10-second rate: 628
    Active flows: 2, Total flows: 10
    Flows exported: 4028, Flows packets exported: 6150
    Flows inactive timed out: 8, Flows active timed out: 4026

    Service Accounting interface: sp-2/1/0, Local interface index: 223
    Flow packets: 0, Flow bytes: 0
    Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
    Active flows: 0, Total flows: 0
    Flows exported: 0, Flows packets exported: 1
    Flows inactive timed out: 0, Flows active timed out: 0
```

show services accounting flow inline-jflow fpc-slot (for IPv4 Flow)

```
user@host> show services accounting flow inline-jflow fpc-slot 5
  Flow information
    FPC Slot: 5
    Flow Packets: 0, Flow Bytes: 0
    Active Flows: 0, Total Flows: 0
    Flows Exported: 0, Flow Packets Exported: 0
    Flows Inactive Timed Out: 0, Flows Active Timed Out: 0
```

show services accounting flow inline-jflow fpc-slot (with IPv4, IPv6, VPLS, and Bridge Configuration)

```
user@host> show services accounting flow inline-jflow fpc-slot 5
Flow information
    FPC Slot: 5
    Flow Packets: 0, Flow Bytes: 0
```

```
    Active Flows: 0, Total Flows: 0
    Flows Exported: 0, Flow Packets Exported: 0
    Flows Inactive Timed Out: 0, Flows Active Timed Out: 0

    IPv4 Flows:
    IPv4 Flow Packets: 0, IPv4 Flow Bytes: 0
    IPv4 Active Flows: 0, IPv4 Total Flows: 0
    IPv4 Flows Exported: 0, IPv4 Flow Packets exported: 0
    IPv4 Flows Inactive Timed Out: 0, IPv4 Flows Active Timed Out: 0

    IPv6 Flows:
    IPv6 Flow Packets: 0, IPv6 Flow Bytes: 0
    IPv6 Active Flows: 0, IPv6 Total Flows: 0
    IPv6 Flows Exported: 0, IPv6 Flow Packets Exported: 0
    IPv6 Flows Inactive Timed Out: 0, IPv6 Flows Active Timed Out: 0

    VPLS Flows:
    VPLS Flow Packets: 0, VPLS Flow Bytes: 0
    VPLS Active Flows: 0, VPLS Total Flows: 0
    VPLS Flows Exported: 0, VPLS Flow Packets Exported: 0
    VPLS Flows Inactive Timed Out: 0, VPLS Flows Active Timed Out: 0

    BRIDGE Flows:
    BRIDGE Flow Packets: 0, BRIDGE Flow Bytes: 0
    BRIDGE Active Flows: 0, BRIDGE Total Flows: 0
    BRIDGE Flows Exported: 0, BRIDGE Flow Packets Exported: 0
    BRIDGE Flows Inactive Timed Out: 0, BRIDGE Flows Active Timed Out: 0
    BRIDGE Flow Insert Count: 0
```

show services accounting flow inline-jflow (MX80 Router with IPv4 and IPv6 Configuration)

```
user@host> show services accounting flow inline-jflow
Flow information
    TFEB Slot: 0
    Flow Packets: 0, Flow Bytes: 0
    Active Flows: 0, Total Flows: 0
    Flows Exported: 0, Flow Packets Exported: 0
    Flows Inactive Timed Out: 0, Flows Active Timed Out: 0

    IPv4 Flows:
    IPv4 Flow Packets: 0, IPv4 Flow Bytes: 0
```

```
    IPv4 Active Flows: 0, IPv4 Total Flows: 0
    IPv4 Flows Exported: 0, IPv4 Flow Packets exported: 0
    IPv4 Flows Inactive Timed Out: 0, IPv4 Flows Active Timed Out: 0

    IPv6 Flows:
    IPv6 Flow Packets: 0, IPv6 Flow Bytes: 0
    IPv6 Active Flows: 0, IPv6 Total Flows: 0
    IPv6 Flows Exported: 0, IPv6 Flow Packets Exported: 0
    IPv6 Flows Inactive Timed Out: 0, IPv6 Flows Active Timed Out: 0
```

show services accounting flow inline-jflow fpc-slot (PTX1000 Router When Both IPv4 and IPv6 Are Configured)

```
user@host> show services accounting flow inline-jflow fpc-slot 0
Flow information
FPC Slot: 0
Flow Packets: 47427946, Flow Bytes: 5217074060
Active Flows: 0, Total Flows: 2
Flows Exported: 194, Flow Packets Exported: 7045
Flows Inactive Timed Out: 2, Flows Active Timed Out: 192

IPv4 Flows:
IPv4 Flow Packets: 47427946, IPv4 Flow Bytes: 5217074060
IPv4 Active Flows: 0, IPv4 Total Flows: 2
IPv4 Flows Exported: 194, IPv4 Flow Packets exported: 7045
IPv4 Flows Inactive Timed Out: 2, IPv4 Flows Active Timed Out: 192

IPv6 Flows:
IPv6 Flow Packets: 0, IPv6 Flow Bytes: 0
IPv6 Active Flows: 0, IPv6 Total Flows: 0
IPv6 Flows Exported: 0, IPv6 Flow Packets Exported: 0
IPv6 Flows Inactive Timed Out: 0, IPv6 Flows Active Timed Out: 0
```

show services accounting flow inline-jflow (Junos OS Evolved 21.2R1 and later, for a PTX10003 router)

```
user@host> show services accounting flow inline-jflow fpc-slot 0
Flow information
  FPC Slot: 0
```

```
  Flow Packets: 0, Flow Bytes: 0
  Active Flows: 0, Total Flows: 0
  Flows Exported: 0, Flow Packets Exported: 0
  Flows Inactive Timed Out: 0, Flows Active Timed Out: 0

  IPv4 Flows:
  IPv4 Flow Packets: 0, IPv4 Flow Bytes: 0
  IPv4 Active Flows: 0, IPv4 Total Flows: 0
  IPv4 Flows Exported: 0
  IPv4 Flows Inactive Timed Out: 0, IPv4 Flows Active Timed Out: 0

  IPv6 Flows:
  IPv6 Flow Packets: 0, IPv6 Flow Bytes: 0
  IPv6 Active Flows: 0, IPv6 Total Flows: 0
  IPv6 Flows Exported: 0
  IPv6 Flows Inactive Timed Out: 0, IPv6 Flows Active Timed Out: 0

  MPLS Flows:
  MPLS Flow Packets: 0, MPLS Flow Bytes: 0
  MPLS Active Flows: 0, MPLS Total Flows: 0
  MPLS Flows Exported: 0
  MPLS Flows Inactive Timed Out: 0, MPLS Flows Active Timed Out: 0
```

show services accounting flow inline-jflow (SRX Series When IPv4 is configured)

```
user@host> show services accounting flow inline-jflow
Flow information
    FPC Slot: 0
    Flow Packets: 462680, Flow Bytes: 45433206
    Active Flows: 34, Total Flows: 61093
    Flows Exported: 138936, Flow Packets Exported: 96649
    Flows Inactive Timed Out: 61083, Flows Active Timed Out: 138936
    Total Flow Insert Count: 0

    IPv4 Flows:
    IPv4 Flow Packets: 462680, IPv4 Flow Bytes: 45433206
    IPv4 Active Flows: 34, IPv4 Total Flows: 61093
    IPv4 Flows Exported: 138936, IPv4 Flow Packets exported: 96649
    IPv4 Flows Inactive Timed Out: 61083, IPv4 Flows Active Timed Out: 138936
    IPv4 Flow Insert Count: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

Junos OS Release 10.0 added the capability to display output from multiple sampling instances.

RELATED DOCUMENTATION

show services accounting status

show services accounting flow-detail

IN THIS SECTION

- Syntax | 478
- Description | 479
- Options | 479
- Additional Information | 480
- Required Privilege Level | 480
- Output Fields | 480
- Sample Output | 483
- Release Information | 485

Syntax

```
show services accounting flow-detail
<detail | extensive | terse>
<filters>
<limit limit-value>
<name (* | all | service-name)>
 <order (bytes | packets)>
```

Description

Display information about the flows being processed by the accounting service.

Options

none — Display information about all flows.

detail | extensive | terse — (Optional) Display the specified level of output.

filters — (Optional) Filter the display output of the currently active flow records. The following filters query actively changing data structures and result in different results for multiple invocations:

- **destination-as**—Display flow records filtered by destination autonomous system information.
- **destination-port**—Display flow records filtered by destination port information.
- **destination-prefix**—Display flow records filtered by destination prefix information.
- **input-snmp-interface-index**—Display flow records filtered by SNMP input interface index information.
- **output-snmp-interface-index**—Display flow records filtered by SNMP output interface index information.
- **proto**—Display flow records filtered by protocol type.
- **source-as**—Display flow records filtered by source autonomous system information.
- **source-port**—Display flow records filtered by source port information.
- **source-prefix**—Display flow records filtered by source prefix information.
- **tos**—Display flow records filtered by type of service classification.

limit *limit-value* — (Optional) Limit the display output to the specified number of flows. The default is no limit.

name (* | all | *service-name*)	(Optional) Display information about the flows being processed. Use a wildcard character, specify all services, or provide a specific services name.
order (bytes | packets)	(Optional) Display the flow with the ordering of the highest number, either by byte count or by packet count.

Additional Information

When no PIC is active, or when no route record has been downloaded from the PIC, this command reports no flows, even though packets are being sampled. This command displays information about two concurrent sessions only. If a third session is attempted, the command pauses with no output until one of the previous sessions is completed.

Required Privilege Level

view

Output Fields

Table 21 on page 480 lists the output fields for the `show services accounting flow-detail` command. Output fields are listed in the approximate order in which they appear.

Table 21: show services accounting flow-detail Output Fields

Field Name	Field Description	Output Level
Service Accounting interface	Name of the service accounting interface.	All levels
Service name	Name of a service that was configured at the `[edit forwarding-options accounting]` hierarchy level. The default display, **(default sampling)**, indicates the service was configured at the `[edit forwarding-options sampling]` hierarchy level.	All levels

Table 21: show services accounting flow-detail Output Fields *(Continued)*

Field Name	Field Description	Output Level
Local interface index	Index counter of the local interface.	All levels
TOS	Type-of-service value from the IP header.	**extensive**
Input SNMP interface index	SNMP index of the interface on which the packet came in.	**extensive**
Output SNMP interface index	SNMP index of the interface on which the packet went out.	**extensive**
Source-AS	Source AS number.	**extensive**
Destination-AS	Destination AS number.	**extensive**
Protocol	Name of the protocol used for the packet flow from the corresponding source address.	All levels
Input interface	Interface on which the packets were received.	All levels
Output interface	Interface on which the packets were transmitted.	All levels
TCP flags	Number of TCP header flags detected in the flow.	**extensive**
Source address	Address where the flow originated.	All levels
Source port	Name of the source port.	All levels
Source prefix length	Source prefix length.	**extensive**

Table 21: show services accounting flow-detail Output Fields *(Continued)*

Field Name	Field Description	Output Level
Destination address	Address where the flow is sent.	All levels
Destination prefix length	Destination prefix length.	**extensive**
Destination port	Name of the destination port.	All levels
Start time	Actual time when the packet in this aggregation was first seen.	**detail extensive**
End time	Actual time when the packet in this aggregation was last seen.	**detail extensive**
Packet count	Number of packets in the aggregation.	All levels
Byte count	Number of bytes in the aggregation.	All levels
Time since last active timeout	Amount of time elapsed since the last active timeout, in the format *hh:mm:ss*.	None specified
Packet count for last active timeout	Number of packets in the aggregation since the last active timeout.	None specified
Byte count for last active timeout	Number of bytes in the aggregation since the last active timeout.	None specified

Sample Output

show services accounting flow-detail

In this sample, the output is split into three sections, with ellipses (...) indicating where the sections are continued.

```
user@host> show services accounting flow-detail
Service Accounting interface: rsp0, Local interface index: 171
Service name: (default sampling)
Interface state: Accounting
Protocol   Input          Source                  Source   Output
           interface      address                   port   interface...
tcp(6)     ge-5/0/1.0     192.0.2.2                    0   ge-5/0/0.0
tcp(6)     ge-5/0/1.0     192.0.2.2                    0   ge-5/0/0.0

Destination          Destination      Packet        Byte  Time since last
address                     port       count       count   active timeout...
198.51.100.149                 0        2660      170240         00:00:58
198.51.100.138                 0        2660      170240         00:00:58

Packet count for       Byte count for
last active timeout    last active timeout
2805                      179520
2805                      179520
```

show services accounting flow-detail limit

In this sample, the output is split into three sections, with ellipses (...) indicating where the sections are continued.

```
user@host> show services accounting flow-detail limit 1
Service Accounting interface: rsp0, Local interface index: 171
Service name: (default sampling)
Interface state: Accounting
Protocol   Input          Source                  Source   Output
           interface      address                 port     interface...
tcp(6)     ge-5/0/1.0     192.0.2.2                    0   ge-5/0/0.0

Destination          Destination      Packet        Byte    Time since last
```

```
address              port             count         count   active timeout...
 198.51.100.149                  0        2158        138112         00:00:47

Packet count for            Byte count for
last active timeout      last active timeout
                     2827                 180928
```

show services accounting flow-detail name extensive

```
user@host> show services accounting flow-detail name cf-2 extensive
Service Accounting interface: mo-0/2/0, Local interface index: 145
Service name: cf-2
  TOS: 0, Protocol: udp(17), TCP flags: 0
  Source address: 10.10.10.1, Source prefix length: 0, Destination address: 203.0.113.20,
Destination prefix length: 0, Source port: 1173, Destination port: 69
  Input SNMP interface index: 65, Output SNMP interface index: 0, Source-AS: 0,
Destination-AS: 0
  Start time: 62425, End time: 635265, Packet count: 165845, Byte count: 9453165
```

show services accounting flow-detail limit order bytes

The output of the following command is displayed over 141 columns, not the standard 80 columns. In this sample, the output is split into three sections, with ellipses (...) indicating where the sections are continued.

```
user@host> show services accounting flow-detail limit 5 order bytes
Service Accounting interface: mo-2/0/0, Local interface index: 356
Service name: (default sampling)
           Input          Source                   Source   Output
Protocol   interface      address                  port     interface...
icmp(1)    ge-2/3/0.0     192.0.2.2                      0 .local.
icmp(1)    ge-2/3/0.0     192.0.2.2                      0 .local.
icmp(1)    ge-2/3/0.0     192.0.2.2                      0 .local.
icmp(1)    ge-2/3/0.0     192.0.2.2                      0 .local.
icmp(1)    ge-2/3/0.0     192.0.2.2                      0 .local.

Destination          Destination       Packet       Byte    Time since last
address              port              count        count   active timeout...
192.168.128.2                   0          16        12148   Not applicable
192.168.144.2                   0          16        15229   Not applicable
```

```
192.168.192.2                  0          16        13296   Not applicable
192.168.16.2                   0          16        13924   Not applicable
192.168.48.2                   0          16        13428   Not applicable

Packet count for          Byte count for
last active timeout       last active timeout
Not applicable            Not applicable
Not applicable            Not applicable
Not applicable            Not applicable
Not applicable            Not applicable
Not applicable            Not applicable
```

show services accounting flow-detail name detail source-port

```
user@host> show services accounting flow-detail name cf-2 detail source-port 1173
Service Accounting interface: mo-0/2/0, Local interface index: 145
Service name: cf-2
  Protocol: udp(17), Source address: 10.10.10.1, Source port: 1173, Destination address:
203.0.113.20, Destination port: 69
  Start time: 62425, End time: 811115, Packet count: 142438, Byte count: 8118966
```

Release Information

Command introduced before Junos OS Release 7.4.

show services accounting memory

IN THIS SECTION

- Syntax | 486
- Description | 486
- Options | 486

- Required Privilege Level | 486
- Output Fields | 486
- Sample Output | 487
- Release Information | 488

Syntax

```
show services accounting  memory
```

Description

Display memory and flow record statistics.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 22 on page 487 lists the output fields for the `show services accounting memory` command. Output fields are listed in the approximate order in which they appear.

Table 22: show services accounting memory Output Fields

Output Field	Output Field Description
Service Accounting interface	Name of the service accounting interface.
Memory Utilization	
Local interface index	Index counter of the local interface.
Allocation count	Number of flow records allocated.
Free count	Number of flow records freed.
Maximum allocated	Maximum number of flow records allocated since the monitoring station booted. This number represents the peak number of flow records allocated at a time.
Allocations per second	Flow records allocated per second during the last statistics interval on the PIC.
Frees per second	Flow records freed per second during the last statistics interval on the PIC.
Total memory used	Total amount of memory currently used (in bytes).
Total memory free	Total amount of memory currently free (in bytes).

Sample Output

show services accounting memory (Monitoring PIC Interface)

```
user@host> show                  services accounting                  memory
Service Accounting interface: mo-2/0/0, Local interface index: 468
  Memory utilization
    Allocation count: 437340, Free count: 433699, Maximum allocated: 6782
    Allocations per second: 3366, Frees per second: 6412
```

```
    Total memory used (in bytes): 133460320,
    Total memory free (in bytes): 133918352
```

show services accounting memory (Service PIC Interface)

```
user@host> show                  services accounting                memory
Service Accounting interface: sp-0/1/0
  Memory utilization
    Allocation count: 1000, Free count: 0
    Allocations per second: 0, Frees per second: 0
    Total memory used (in bytes): 218158272
    Total memory free (in bytes): 587147696
```

```
Service Accounting interface: sp-1/0/0
  Memory utilization
    Allocation count: 1000, Free count: 0
    Allocations per second: 0, Frees per second: 0
    Total memory used (in bytes): 218157592
    Total memory free (in bytes): 587148376
```

Release Information

Command introduced before Junos OS Release 7.4.

show services accounting packet-size-distribution

IN THIS SECTION

- Syntax | 489
- Description | 489

- Options | 489
- Required Privilege Level | 489
- Output Fields | 490
- Sample Output | 490
- Release Information | 491

Syntax

```
show services accounting packet-size-distribution
<name (* | all | service-name)>
```

Description

Display a packet size distribution histogram.

Options

none	Display a packet size distribution histogram of all accounting services.
name (* | all | ***service-name*****)**	(Optional) Display a packet size distribution histogram. Use a wildcard character, specify all services, or provide a specific services name.

Required Privilege Level

view

Output Fields

Table 23 on page 490 lists the output fields for the `show services accounting packet-size-distribution` command. Output fields are listed in the approximate order in which they appear.

Table 23: show services accounting packet-size-distribution Output Fields

Field Name	Field Description
Service Accounting interface	Name of the service accounting interface.
Service name	Name of a service that was configured at the [**edit-forwarding-options accounting**] hierarchy level. The default display, **(default sampling)**, indicates the service was configured at the [**edit-forwarding-options sampling-level**] hierarchy level.
Local interface index	Index counter of the local interface.
Range start	Smallest packet length (in bytes) to count.
Range end	Largest packet length (in bytes) to count.
Number of packets	Count of packets detected in the size between **Range start** and **Range end**.
Percentage packets	Percentage of the total number of packets that are in this size range.

Sample Output

show services accounting packet-size-distribution name

```
user@host> show services accounting packet-size-distribution name test3
Service Accounting interface: mo-0/2/0, Local interface index: 163
Service name: test3
Range start     Range end     Number of packets     Percentage packets
           32            64                  2924                    100
```

Release Information

Command introduced before Junos OS Release 7.4.

show services accounting status

IN THIS SECTION

- Syntax | 491
- Description | 491
- Options | 492
- Required Privilege Level | 492
- Output Fields | 492
- Sample Output | 494
- Sample Output | 494
- Release Information | 496

Syntax

```
show services accounting  status
<inline-jflow fpc-slot slot-number | name (* | all | service-name)>
```

Description

Display available Physical Interface Cards (PICs) for accounting services.

Options

none	Display available PICs for all accounting services.
inline-jflow fpc-slot *slot-number*	(Optional) Display inline flow accounting status for the specified FPC. For a two-member MX Series Virtual Chassis or EX9200 Virtual Chassis, the primary router or switch uses FPC slot numbers 0 through 11 with no offset; the backup router or switch uses FPC slot numbers 12 through 23, with an offset of 12.
name (* | all | *service-name*)	(Optional) Display available PICs. Use a wildcard character, specify all services, or provide a specific services name.

Required Privilege Level

view

Output Fields

Table 24 on page 492 lists the output fields for the `show services accounting status` command. Output fields are listed in the approximate order in which they appear.

Table 24: show services accounting status Output Fields

Field	Field Description
`Service Accounting interface`	Name of the service accounting interface.
`Service name`	Name of a service that was configured at the [`edit-forwarding-options accounting`] hierarchy level. The default display,(`default sampling`), indicates the service was configured at the [`edit-forwarding-options sampling-level`] hierarchy level.
`FPC Slot`	Slot number of the FPC for which the flow information is displayed.

Table 24: show services accounting status Output Fields *(Continued)*

Field	Field Description
`Local interface index`	Index counter of the local interface.
`Interface state`	Accounting state of the passive monitoring interface. • `Accounting`—PIC is actively accounting. • `Disabled`—PIC has been disabled from the CLI. • `Not accounting`—PIC is up but not accounting. This can happen while the PIC is coming online, or when the PIC is up but has no logical unit configured under the physical interface. • Unknown
`Group index`	Integer that represents the monitoring group of which the PIC is a member. `Group index` is a mapping from the group name to an index. It is not related to the number of monitoring groups.
`Export interval (in seconds)`	Configured export interval for cflowd records, in seconds.
`Export format`	Configured export format.
`Protocol`	Protocol the PIC is configured to monitor.
`Engine type`	Configured engine type that is inserted in output cflowd packets.
`Engine ID`	Configured engine ID that is inserted in output cflowd packets.
`Route Record Count`	Number of routes recorded.
`AS Record Count`	Number of autonomous systems recorded.

Table 24: show services accounting status Output Fields *(Continued)*

Field	Field Description
Route Records Set	Status of route recording; whether routes are recorded or not.
Configuration Set	Status of monitoring configuration; whether monitoring configuration is set or not.

Sample Output

show services accounting status name (Monitoring PIC Interface)

```
user@host> show services accounting status name count1
Service Accounting interface: mo-2/0/0, Local interface index: 468
Service name: count1
Interface state: Accounting
  Group index: 0
  Export interval (in seconds): 60, Export format: cflowd v8
  Protocol: IPv4, Engine type: 55, Engine ID: 5
```

Sample Output

show services accounting status name (Service PIC Interface)

```
user@host> show services accounting status name
Service Accounting interface: sp-0/1/0
Interface state: Accounting
  Export format: 9, Route record count: 0
  IFL to SNMP index count: 7, AS count: 0
  Configuration set: Yes, Route record set: No, IFL SNMP map set: Yes
```

```
Service Accounting interface: sp-1/0/0
Interface state: Accounting
  Export format: 9, Route record count: 33
  IFL to SNMP index count: 7, AS count: 1
  Configuration set: Yes, Route record set: Yes, IFL SNMP map set: Yes
```

show services accounting status inline-jflow fpc-slot (When IPv4, IPv6 and Bridge Family Are Configured)

```
user@host> show services accounting status inline-jflow fpc-slot 0
FPC Slot: 0
     IPV4 export format: Version-IPFIX, IPV6 export format: Not set
     BRIDGE export format: Version-IPFIX, MPLS export format: Version-IPFIX
     IPv4 Route Record Count: 31, IPv6 Route Record Count: 0, MPLS Route Record Count: 13
     Route Record Count: 44, AS Record Count: 1
     Route-Records Set: Yes, Config Set: Yes
     Service Status: PFE-0: Steady PFE-1: Steady
     Using Extended Flow Memory?: PFE-0: No PFE-1: No
     Flex Flow Sizing ENABLED?: PFE-0: No PFE-1: No
     IPv4 MAX FLOW Count: 1024, IPv6 MAX FLOW Count: 512
     BRIDGE MAX FLOW Count: 1024, MPLS MAX FLOW Count: 1024
     MAX Flow Table size: 15
```

show services accounting status inline-jflow (MX80 Router When Both IPv4 and IPv6)

```
user@host> show services accounting status inline-jflow

Status information
          TFEB Slot: 0
    Export format: IP-FIX
    IPv4 Route Record Count: 6, IPv6 Route Record Count: 8
    Route Record Count: 14, AS Record Count: 1
    Route-Records Set: Yes, Config Set: Yes
```

show services accounting status inline-jflow fpc-slot (PTX1000 Router When Both IPv4 and IPv6 Are Configured)

```
user@host> show services accounting status inline-jflow fpc-slot 0
Status information
FPC Slot: 0
IPV4 export format: Version-IPFIX, IPV6 export format: Version-IPFIX
MPLS export format: Not set
IPv4 Route Record Count: 23, IPv6 Route Record Count: 3, MPLS Route Record Count: 0
Route Record Count: 26, AS Record Count: 1
Route-Records Set: Yes, Config Set: Yes
```

show services accounting status inline-jflow (SRX Series Devices When Both IPv4 and IPv6 Are Configured)

```
user@host> show services accounting status inline-jflow
Status information
    FPC Slot: 0
    IPV4 export format: Version9, IPV6 export format: Version9
    BRIDGE export format: Not set, MPLS export format: Not set
    IPv4 Route Record Count: 24, IPv6 Route Record Count: 0, MPLS Route Record Count: 0
    Route Record Count: 24, AS Record Count: 1
    Route-Records Set: Yes, Config Set: Yes
    Service Status: PFE-0: Steady
    Using Extended Flow Memory?: PFE-0: No
    Flex Flow Sizing ENABLED?: PFE-0: No
    IPv4 MAX FLOW Count: 0, IPv6 MAX FLOW Count: 0
    BRIDGE MAX FLOW Count: 0, MPLS MAX FLOW Count: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

show services accounting flow

Inline Flow Monitoring for Virtual Chassis Overview

show services accounting usage

IN THIS SECTION

- Syntax | 497
- Description | 497
- Options | 497
- Additional Information | 498
- Required Privilege Level | 498
- Output Fields | 498
- Sample Output | 499
- Release Information | 500

Syntax

```
show services accounting usage
<name service-name>
```

Description

Display the CPU usage of PIC used for active flow monitoring.

Options

none	Display CPU usage for all service names.

name *service-name* (Optional) Display CPU usage for the specified service name.

Additional Information

When no route record has been downloaded from the PIC, this command reports no flows, even though packets are being sampled.

Required Privilege Level

view

Output Fields

Table 25 on page 498 lists the output fields for the `show services accounting usage` command. Output fields are listed in the approximate order in which they appear.

Table 25: show services accounting usage Output Fields

Output Field	Output Field Description
Service Accounting interface	Name of the service accounting interface.
Service name	Name of a service that was configured at the [**edit-forwarding-options accounting**] hierarchy level. The default display, **(default sampling)**, indicates the service was configured at the [**edit-forwarding-options sampling-level**] hierarchy level.
Local interface index	Index counter of the local interface.
Uptime	Time that the PIC has been operational (in milliseconds).
Interrupt time	Total time that the PIC has spent processing packets since the last PIC reset (in microseconds).

Table 25: show services accounting usage Output Fields *(Continued)*

Output Field	Output Field Description
Load (5 second)	CPU load on the PIC, averaged more than 5 seconds. The number is a percentage obtained by dividing the time spent on active tasks by the total elapsed time.
Load (1 minute)	CPU load on the PIC, averaged more than 1 minute. The number is a percentage obtained by dividing the time spent on active tasks by the total elapsed time.

Sample Output

show services accounting usage (Monitoring PIC Interface)

```
user@host> show services accounting usage
Service Accounting interface: mo-1/1/0, Local interface index: 15
Service name: (default sampling)
  CPU utilization
    Uptime: 600413856 milliseconds, Interrupt time: 2403 microseconds
    Load (5 second): 43%, Load (1 minute): 24%
```

show services accounting usage (Service PIC Interface)

```
user@host> show services accounting usage
Service Accounting interface: sp-0/1/0
Service name: (default sampling)
  CPU utilization
    Uptime: 7853940 milliseconds, Interrupt time: 0 microseconds
    Load (5 second): 2%, Load (1 minute): 0%



Service Accounting interface: sp-0/1/0
Service name: (default sampling)
  CPU utilization
```

```
    Uptime: 331160 milliseconds, Interrupt time: 0 microseconds
    Load (5 second): 2%, Load (1 minute): 0%
```

Release Information

Command introduced before Junos OS Release 7.4.

show services flow-collector file interface

IN THIS SECTION

- Syntax | **500**
- Description | **501**
- Options | **501**
- Additional Information | **501**
- Required Privilege Level | **501**
- Output Fields | **501**
- Sample Output | **503**
- Release Information | **504**

Syntax

```
show services flow-collector file interface (all | cp-fpc/pic/port)
<detail | extensive | terse>
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display information about flow collector files.

Options

none	Display file information for all configured flow collector interfaces.
all | cp-*fpc/pic/port*	Display file information for all configured flow collector interfaces or for the specified interface.
detail | extensive | terse	(Optional) Display the specified level of output.

Additional Information

No entries are displayed for files that have been successfully transferred.

Required Privilege Level

view

Output Fields

Table 26 on page 501 lists the output fields for the `show services flow-collector file interface` command. Output fields are listed in the approximate order in which they appear.

Table 26: show services flow-collector file interface Output Fields

Output Field	Output Field Description	Level of Output
Filename	Name of the file created on the flow collector interface.	All levels

Table 26: show services flow-collector file interface Output Fields *(Continued)*

Output Field	Output Field Description	Level of Output
Flows	Total number of collector flows for which records are present in the file.	none specified

Throughput	Throughput statistics:
	• **Flow records**—Number of flow records in the file.
	  • **per second**—Average number of flow records per second.
	  • **peak per second**—Peak number of flow records per second.
	• **Uncompressed bytes**—Total file size before compression.
	  • **per second**—Average number of uncompressed bytes per second.
	  • **peak per second**—Peak number of uncompressed bytes per second.
	• **Compressed bytes**—Total file size after compression.
	  • **per second**—Average number of compressed bytes per second.
	  • **peak per second**—Peak number of compressed bytes per second.
		extensive

Table 26: show services flow-collector file interface Output Fields *(Continued)*

Output Field	Output Field Description	Level of Output

Status	File statistics:
	• **Compressed blocks**—(**extensive** output only) Data blocks in the file that have been compressed. The file is exported only when the compressed block count and block count become the same.
	• **Block count**—(**extensive** output only) Total number of data blocks in the file.
	• **State**—Processing state of the file.
	  • **Active**—The flow collector interface is writing to the file.
	  • **Export 1**—File export is in progress to the primary server.
	  • **Export 2**—File export is in progress to the secondary server.
	  • **Wait**—File is pending export.
	  • **Transfer attempts 0**.—Number of attempts made to transfer the file. If the file is successfully transferred in the first attempt, this field is 0.
		All levels

Sample Output

show services flow-collector file interface extensive

```
user@host> show services flow-collector file interface cp-3/2/0 extensive
Filename: cFlowd-py69Ni69-0-20031112_014301-so_3_0_0_0.bcp.bi.gz
  Throughput:
    Flow records: 188365, per second: 238, peak per second: 287
    Uncompressed bytes: 21267756, per second: 27007, peak per second: 32526
    Compressed bytes: 2965643, per second: 0, peak per second: 22999
  Status:
    Compressed blocks: 156, Block count: 156
    State: Active, Transfer attempts: 0
```

Release Information

Command introduced before Junos OS Release 7.4.

show services flow-collector input interface

IN THIS SECTION

- Syntax | 504
- Description | 504
- Options | 505
- Required Privilege Level | 505
- Output Fields | 505
- Sample Output | 506
- Release Information | 507

Syntax

```
show services flow-collector input interface (all | cp-fpc/pic/port)
<detail | extensive | terse>
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display the number of packets received by collector interfaces from monitoring interfaces.

Options

none	Display packets received by all configured flow collector interfaces.
all | cp-*fpc/pic/port*	Display packets received by all configured flow collector interfaces or by the specified interface.
detail | extensive | terse	(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 27 on page 505 lists the output fields for the `show services flow-collector input interface` command. Output fields are listed in the approximate order in which they appear.

Table 27: show services flow-collector input interface Output Fields

Output Field	Output Field Description
Interface	Name of the monitoring interface.
Packets	Number of packets traveling from the monitoring interface to the flow collector interface.
Bytes	Number of bytes traveling from the monitoring interface to the flow collector interface.

Sample Output

show services flow-collector input interface

```
user@host> show services flow-collector input interface cp-3/2/0
Interface                      Packets        Bytes
mo-3/0/0.0                       21706     32328568
mo-3/1/0.0                       21706     32329096
```

show services flow-collector input interface all

```
user@host> show services flow-collector input interface all
Flow collector interface: cp-6/1/0
Interface state: Collecting flows
Interface                      Packets        Bytes
mo-3/0/0.0                         274       416232
mo-3/3/0.0                         274       416184
mo-1/0/0.0                         274       416232
mo-1/1/0.0                         274       416232
mo-1/2/0.0                         274       416232
mo-1/3/0.0                         274       416232
mo-3/1/0.0                         274       416232
mo-4/0/0.0                         274       416232
mo-4/1/0.0                         274       416232
mo-4/2/0.0                         274       416184
mo-4/3/0.0                         274       416232
mo-5/0/0.0                         274       416232
mo-5/1/0.0                         274       416232
mo-5/2/0.0                         274       416232
mo-5/3/0.0                         274       416232
mo-6/0/0.0                         274       416232

Flow collector interface: cp-6/3/0
Interface state: Collecting flows
```

Release Information

Command introduced before Junos OS Release 7.4.

show services flow-collector interface

IN THIS SECTION

- Syntax | 507
- Description | 507
- Options | 508
- Required Privilege Level | 508
- Output Fields | 508
- Sample Output | 514
- Release Information | 518

Syntax

```
show services flow-collector interface (all | cp-fpc/pic/port)
<detail | extensive | terse>
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display overall statistics for the flow collector application.

Options

none	Display statistics for flow collector applications on all interfaces.
all | cp-*fpc/pic/port*	Display statistics for flow collector applications on all interfaces or for the specified interface.
detail | extensive | terse	(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

Table 28 on page 508 lists the output fields for the `show services flow-collector interface` command. Output fields are listed in the approximate order in which they appear.

Table 28: show services flow-collector interface Output Fields

Output Field	Output Field Description	Level of Output
Flow collector interface	Name of the flow collector interface.	All levels
Interface state	Collecting flow state for the interface.	All levels
Packets	Total number of packets received.	none specified
Flows Uncompressed Bytes	Total uncompressed data size for all files created on this PIC.	none specified
Compressed Bytes	Total compressed data size for all files created on this PIC.	none specified

Table 28: show services flow-collector interface Output Fields *(Continued)*

Output Field	Output Field Description	Level of Output
FTP bytes	Total number of bytes transferred to the FTP server, including those dropped during transfer.	none specified
FTP files	Total number of FTP transfers attempted by the server.	none specified
Memory	Bytes used on the PIC and bytes free.	**detail extensive**

Input	Incoming flow collector packet statistics:
	• **Packets**—Number of packets received on the unit.
	  • **per second**—Average number of packets per second.
	  • **peak per second**—Peak number of packets per second.
	• **Bytes**—Number of bytes received on the unit.
	  • **per second**—Average number of bytes per second.
	  • **peak per second**—Peak number of bytes per second.
	• **Flow records processed**—Number of records in the flow collector packets that were processed by the flow-collector interface.
	  • **per second**—Average number of flow records processed per second.
	  • **peak per second**—Peak number of flow records per second.
		detail extensive

Table 28: show services flow-collector interface Output Fields *(Continued)*

Output Field	Output Field Description	Level of Output

Allocation	Data block statistics:
	• **Blocks allocated**—Total number of data blocks (containing flow records) allocated to the files created on this PIC.
	  • **per second**—Average number of blocks allocated per second.
	  • **peak per second**—Peak number of blocks allocated per second.
	• **Blocks freed**—Total number of data blocks freed.
	  • per second—Average number of blocks freed per second.
	  • peak per second—Peak number of blocks freed per second.
	• **Blocks unavailable**—Total number of data block requests denied, typically because of a memory shortage.
	  • per second—Average number of blocks unavailable per second.
	  • peak per second—Peak number of blocks unavailable per second.
		extensive

Files	File statistics, incremented since the PIC last booted: • **Files created**—Total number of files created on this PIC. • **Files exported**— Number of files successfully created and exported. • **Files destroyed**—(**extensive** output only) Number of files successfully exported and files dropped by the flow collection interface.	**detail extensive**

Table 28: show services flow-collector interface Output Fields *(Continued)*

Output Field	Output Field Description	Level of Output

Throughput	Throughput statistics:
	• **Uncompressed bytes**—Total uncompressed data size for all files created on this PIC.
	  • **per second**—Average number of uncompressed bytes per second.
	  • **peak per second**—Peak number of uncompressed bytes per second.
	• **Compressed bytes**—Total compressed data size for all files created on this PIC.
	  • **per second**—Average number of compressed bytes per second.
	  • **peak per second**—Peak number of compressed bytes per second.
		detail extensive

Packet drops	Number of packets dropped for the following causes: • **No memory**—Packets dropped because of insufficient memory. • **Not IP**—Packets dropped because they are not IP packets. • **Not IPv4**—Packets dropped because they are not IP version 4 packets. • **Too small**—Packets dropped because each packet was smaller than the size reported in its header. • **Fragments**—Packets dropped because of fragmentation. Fragments are not reassembled. • **ICMP**—Packets dropped because they are not ICMP packets. • **TCP**—Packets dropped because they are not TCP packets. • **Unknown**—Packets dropped because of undetermined causes. • **Not Junos flow**—Packets dropped because they are not interpreted by Junos OS. Junos OS interprets only IPv4, UDP cflowd version 5 packets.	**extensive**

Table 28: show services flow-collector interface Output Fields *(Continued)*

Output Field	Output Field Description	Level of Output
File transfer	File transfer statistics: • **FTP bytes**—Total number of bytes transferred to the FTP server, including those dropped during transfer. • **FTP files**—Total number of FTP transfers attempted by the server. • **FTP failure**—Total number of FTP failures encountered by the server.	**detail extensive**
Flow collector interface	Physical interface acting as a flow collector.	**detail**

Table 28: show services flow-collector interface Output Fields *(Continued)*

Output Field	Output Field Description	Level of Output

Export channel	Export channel **0** is unit 0. Export channel **1** is unit 1. Flow receive channel is unit 2. Server status statistics are the following:
	• **Current server Primary** or **Secondary**—Current FTP server being used. Value is
	• **Primary server state**—State of the server:
	  • **OK**—Server is operating without problems.
	  • **FTP error**—Server encountered an FTP protocol error while sending files.
	  • **Network error**—Flow-collector interface has errors when contacting the primary FTP server.
	  • **Unknown**—First file transfer has not been sent to the primary server.
	• **Secondary server state**—State of the server:
	  • **OK**—Server is operating without errors.
	  • **FTP error**—Server encountered an FTP protocol error while sending files.
	  • **Network error**—Flow-collector interface has errors when contacting the secondary FTP server.
	  • **Unknown**—First file transfer has not been sent to the secondary server.
	• **Not configured**—Secondary server is not configured.
		detail extensive

Sample Output

show services flow-collector interface all detail

```
user@host> show services flow-collector interface all detail
Flow collector interface: cp-6/1/0
Interface state: Collecting flows
Memory:
    Used: 51452732, Free: 440329088
Input:
    Packets: 4384, per second: 0, peak per second: 156
    Bytes: 6659616, per second: 0, peak per second: 249695
    Flow records processed: 131070, per second: 0, peak per second: 4914
Files:
    Files created: 1, per second: 0, peak per second: 0
    Files exported: 1, per second: 0, peak per second: 0
Throughput:
    Uncompressed bytes: 13742307, per second: 0, peak per second: 593564
    Compressed bytes: 3786177, per second: 0, peak per second: 162826
File Transfer:
    FTP bytes: 3786247, per second: 0, peak per second: 378620
    FTP files: 1, per second: 0, peak per second: 0
    FTP failure: 0
Export channel: 0
    Current server: Primary
    Primary server state: OK, Secondary server state: OK
Export channel: 1
    Current server: Primary
    Primary server state: Unknown, Secondary server state: OK

Flow collector interface: cp-6/3/0
Interface state: Collecting flows
Memory:
    Used: 51452732, Free: 440329088
Input:
    Packets: 0, per second: 0, peak per second: 0
    Bytes: 0, per second: 0, peak per second: 0
    Flow records processed: 0, per second: 0, peak per second: 0
Files:
    Files created: 0, per second: 0, peak per second: 0
    Files exported: 0, per second: 0, peak per second: 0
```

```
Throughput:
    Uncompressed bytes: 0, per second: 0, peak per second: 0
    Compressed bytes: 0, per second: 0, peak per second: 0
File Transfer:
    FTP bytes: 70, per second: 0, peak per second: 6
    FTP files: 0, per second: 0, peak per second: 0
    FTP failure: 0
Export channel: 0
    Current server: Primary
    Primary server state: Unknown, Secondary server state: OK
Export channel: 1
    Current server: Primary
    Primary server state: Unknown, Secondary server state: OK
```

show services flow-collector interface all extensive

```
user@host> show services flow-collector interface all extensive
Flow collector interface: cp-6/1/0
Interface state: Collecting flows
Memory:
    Used: 51452732, Free: 440329088
Input:
    Packets: 4384, per second: 0, peak per second: 156
    Bytes: 6659616, per second: 0, peak per second: 249695
    Flow records processed: 131070, per second: 0, peak per second: 4914
Allocation:
    Blocks allocated: 108, per second: 0, peak per second: 0
    Blocks freed: 108, per second: 0, peak per second: 10
    Blocks unavailable: 0, per second: 0, peak per second: 0
Files:
    Files created: 1, per second: 0, peak per second: 0
    Files exported: 1, per second: 0, peak per second: 0
    Files destroyed: 1, per second: 0, peak per second: 0
Throughput:
    Uncompressed bytes: 13742307, per second: 0, peak per second: 593564
    Compressed bytes: 3786177, per second: 0, peak per second: 162826
Packet drops:
    No memory: 0, Not IP: 0
    Not IPv4: 0, Too small: 0
    Fragments: 0, ICMP: 0
    TCP: 0, Unknown: 0
```

```
    Not JUNOS flow: 0
File Transfer:
    FTP bytes: 3786247, per second: 0, peak per second: 378620
    FTP files: 1, per second: 0, peak per second: 0
    FTP failure: 0
Export channel: 0
    Current server: Primary
    Primary server state: OK, Secondary server state: OK
Export channel: 1
    Current server: Primary
    Primary server state: Unknown, Secondary server state: OK

Flow collector interface: cp-6/3/0
Interface state: Collecting flows
Memory:
    Used: 51452732, Free: 440329088
Input:
    Packets: 0, per second: 0, peak per second: 0
    Bytes: 0, per second: 0, peak per second: 0
    Flow records processed: 0, per second: 0, peak per second: 0
Allocation:
    Blocks allocated: 0, per second: 0, peak per second: 0
    Blocks freed: 0, per second: 0, peak per second: 0
    Blocks unavailable: 0, per second: 0, peak per second: 0
Files:
    Files created: 0, per second: 0, peak per second: 0
    Files exported: 0, per second: 0, peak per second: 0
    Files destroyed: 0, per second: 0, peak per second: 0
Throughput:
    Uncompressed bytes: 0, per second: 0, peak per second: 0
    Compressed bytes: 0, per second: 0, peak per second: 0
Packet drops:
    No memory: 0, Not IP: 0
    Not IPv4: 0, Too small: 0
    Fragments: 0, ICMP: 0
    TCP: 0, Unknown: 0
    Not JUNOS flow: 0
File Transfer:
    FTP bytes: 70, per second: 0, peak per second: 6
    FTP files: 0, per second: 0, peak per second: 0
    FTP failure: 0
Export channel: 0
    Current server: Primary
```

```
    Primary server state: Unknown, Secondary server state: OK
Export channel: 1
    Current server: Primary
    Primary server state: Unknown, Secondary server state: OK
```

show services flow-collector interface all terse

```
user@host> show services flow-collector interface all terse
Flow collector interface: cp-6/1/0
Interface state: Collecting flows
  Packets     Bytes     Flows Uncompressed  Compressed     FTP bytes FTP files
                                    Bytes       Bytes
     4384   6659616    131070     13742307     3786177       3786247         1

Flow collector interface: cp-6/3/0
Interface state: Collecting flows
  Packets     Bytes     Flows Uncompressed  Compressed     FTP bytes FTP files
                                    Bytes       Bytes
        0         0         0            0           0            70         0
```

show services flow-collector interface extensive

```
user@host> show services flow-collector interface cp-5/2/0 extensive
Flow collector interface: cp-5/2/0
Interface state: Collecting flows
Memory:
     Used: 458311860, Free: 40810008
Input:
     Packets: 922629, per second: 2069, peak per second: 3266
     Bytes: 1376559252, per second: 3096940, peak per second: 4880051
     Flow records processed: 25764957, per second: 42564, peak per second: 98124
Allocation:
     Blocks allocated: 20862, per second: 31, peak per second: 72
     Blocks freed: 17161, per second: 40, peak per second: 202
     Blocks unavailable: 58786, per second: 652, peak per second: 1120
Files:
     Files created: 52, per second: 0, peak per second: 0
     Files exported: 42, per second: 0, peak per second: 0
     Files destroyed: 42, per second: 0, peak per second: 0
Throughput:
```

```
     Uncompressed bytes: 2592070401, per second: 7297307,
     peak per second: 8630023
     Compressed bytes: 659600068, per second: 1858458, peak per second: 2198471
Packet drops:
     No memory: 58786, Not IP: 0
     Not IPv4: 0, Too small: 0
     Fragments: 0, ICMP: 0
     TCP: 0, Unknown: 0
     Not JUNOS flow: 0
File Transfer:
     FTP bytes: 585981447, per second: 1313320, peak per second: 4857798
     FTP files: 48, per second: 0, peak per second: 0
     FTP failure: 8
Export channel: 0
     Current server: Primary
     Primary server state: FTP error, Secondary server state: Not configured Export channel: 1
     Current server: Primary
     Primary server state: OK, Secondary server state: Not configured
```

Release Information

Command introduced before Junos OS Release 7.4.